Calhoun: The NPS Institutional Archive

DSpace Repository

Theses and Dissertations | 1. Thesis and Dissertation Collection, all items

1995-09

Migrating to ATM through multilayer frame switched networking

Kelly, James W.

Monterey, California. Naval Postgraduate School

http://hdl.handle.net/10945/39140

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun

http://www.nps.edu/library

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA

THESIS

MIGRATING TO ATM THROUGH MULTILAYER FRAME SWITCHED NETWORKING

by

James W. Kelly

19960207 008

September, 1995

Thesis Advisor: Myung Suh
Associate Advisor: Rex Buddenberg

Approved for public release; distribution is unlimited.

DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE September 1995	3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE MIGRATING TO ATM THROUGH MULTILAYER FRAME SWITCHED NETWORKING

5. FUNDING NUMBERS

6. AUTHOR(S) Kelly, James W.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(maximum 200 words)*

Asynchronous Transfer Mode (ATM) technology has received an unprecedented level of worldwide acceptance as the emerging standard for information communication. Transitioning to new information technologies routinely exposes organizations to new difficulties. Well-considered migration strategies will facilitate the implementation of ATM technology.

This study performs an architectural analysis of using frame switched enterprise networking as a strategic step in migrating from conventional networks to ATM networks. Application of the strategy to a multi-vendor, multi-protocol internet, the U.S. Department of Transportation Intermodal Data Network, is considered.

14. SUBJECT TERMS ATM; Frame Switching; Cell Relay; Network Architecture; Network Topology; Enterprise Network; Wide Area Network; Local Area Network			15. NUMBER OF PAGES 79
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18 298-102

Approved for public release; distribution is unlimited.

MIGRATING TO ATM THROUGH MULTILAYER FRAME SWITCHED NETWORKING

James W. Kelly
Lieutenant Commander, United States Coast Guard
B.S., United States Coast Guard Academy, 1978

Submitted in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
September 1995

Author: [redacted]
James W. Kelly

Approved by: [redacted]
Myung Suh, Thesis Advisor

[redacted]
Rex Buddenberg, Associate Advisor

[redacted]
Reuben T. Harris, Chairman, Department of Systems Management

ABSTRACT

Asynchronous Transfer Mode (ATM) technology has received an unprecedented level of worldwide acceptance as the emerging standard for information communication. Transitioning to new information technologies routinely exposes organizations to new difficulties. Well-considered migration strategies will facilitate the implementation of ATM technology.

This study performs an architectural analysis of using frame switched enterprise networking as a strategic step in migrating from conventional networks to ATM networks. Application of the strategy to a multi-vendor, multi-protocol internet, the U.S. Department of Transportation Intermodal Data Network, is considered.

TABLE OF CONTENTS

I. INTRODUCTION 1

A. THE EVOLUTIONARY STATE OF NETWORKING 1

B. OBJECTIVE 2

C. METHODOLOGY 2

D. ORGANIZATION 3

II. ANALYSIS OF FRAME SWITCHING 5

A. CONTEXT 5

B. SWITCHING NETWORKS 8

1. Frame Switch Locations 9

2. Characteristics of Frame Switching 10

3. Additional Applications for Switches 11

C. ARCHITECTURES OF SWITCHED NETWORKS 11

1. Distributed Backbone 11

2. Distributed Backbone With Routing 12

3. Distributed Backbone With Switching and Routing 13

4. Cross-Architectural Issues 15

D. COST AND VALUE OF SWITCHING 16

E. ANALYSIS OF SWITCH ARCHITECTURES 17

1. Crossbar Switches 18

2. Shared Memory Switch Architecture 20

3. Shared Bus Switch Architecture 21

4. Cell Bus Switch Architecture 22

F. VIRTUAL LANs: THE STRATEGIC DRIVER 23

III. ANALYSIS OF ATM CELL SWITCHING 27

A. ATM IN CONTEXT 27

B. VIRTUAL CHANNELS AND VIRTUAL PATHS 28
1. ATM Clocking 28
2. Aggregation of Connections 29
3. Permanent and Switched Virtual Connections 30
C. ATM EQUIPMENT 30
1. Characteristics of ATM Switches 30
2. Characteristics of ATM Adapters 32
D. ANALYSIS OF ATM SWITCH ATTRIBUTES 33
1. Edge Switches 33
2. Workgroup Switches 33
3. Enterprise Switches 34
E. SWITCH SELECTION CRITERIA 34
F. CONNECTION MANAGEMENT 36
G. MULTILAYER SWITCHING CHALLENGES 37

IV. APPLICATION OF THE MIGRATION STRATEGY 41
A. DOT NETWORKING ENVIRONMENT 41
1. The United States Department of Transportation 41
2. Decentralized Control 42
3. Heterogeneous Networks 43
B. NETWORK DIAGRAMS 43
C. IDN FDDI MAIN RING CONFIGURATION 44
D. CHARACTERISTICS OF THE NETBUILDER II 46
E. MODELING IDN TRAFFIC 46
F. MODELING RESULTS 48
G. LINK DELAYS AND UTILIZATION 50
H. TRADEOFFS IN IMPLEMENTING SWITCHING IN THE IDN 51
I. SWITCHING ARCHITECTURES APPLIED TO THE IDN 53
1. Applying Switching To Server Connections 53

2. Applying Switching To Subnetwork LANs 56
J. RECOMMENDATIONS FOR MIGRATING TO SWITCHING NETWORKS 56
K. CONCLUSION 57

LIST OF REFERENCES 59

BIBLIOGRAPHY 63

INITIAL DISTRIBUTION LIST 67

I. INTRODUCTION

A. THE EVOLUTIONARY STATE OF NETWORKING

Asynchronous Transfer Mode (ATM) is a new telecommunications technology which provides a large amount of digital bandwidth and a common format for services with differing bandwidth and Quality of Service requirements.

> It is clear that Asynchronous Transfer Mode (ATM) technology will play a central role in the evolution of current workgroup, campus, and enterprise networks. ATM delivers important advantages over existing LAN [Local Area Network] and WAN [Wide Area Network] technologies, including the promise of scalable bandwidths at unprecedented price and performance points and Quality of Service (QoS) guarantees, which facilitate new classes of applications such as multimedia.
>
> These benefits, however, come at a price. Contrary to common misconceptions, ATM is a very complex technology, perhaps the most complex ever developed by the networking industry. (Alles1)

Decreasing marginal costs for bandwidth are making the migration to new technologies attractive. Eric Benhamou, 3Com's chairman and CEO, said that the cost of obtaining a 10-fold performance increase just by migrating to Fast Ethernet will only be double the cost of current networks. Benhamou predicts that Fast Ethernet "...will become the de facto standard for desktops, because the incremental cost will be minimal." (WSJ) Adding ATM switches to Fast Ethernet LANs can provide yet another ten-fold increase in bandwidth.

The $8 billion dollar internetworking equipment business' growth rate has been phenomenal. Cisco Systems' 1994 sales of hubs and routers were double its 1993 level and 18 times its 1990 sales figure. (WSJ) FORE Systems, a manufacturer of ATM switches, is currently growing faster than Cisco ever grew. (Lippis)

The combination of faster data rates and new frame switches will deliver enormous power to desktops and drive industry growth to as much as 50% yearly, according to Thomas Pincince, a senior analyst at Forrester Research Inc. Pincince believes that the demand is already strong for such power, citing applications such as team projects,

document imaging, and Web browsing. The industry appears ready to deliver the new infrastructure to serve that customer need, according to Pincince. (WSJ)

B. OBJECTIVE

The primary purpose of this thesis is to provide structure to a study of ATM and frame switching. The secondary purpose is to examine the strategic use of frame switching as a risk-mitigating factor in the migration from conventional to ATM networks. A key basis for this study's recommended migration strategy is that *switched LANs provide a migratory path to ATM* (Morency1). Toward this end the application of the strategy to a multi-vendor, multi-protocol network, the U.S. Department of Transportation's Intermodal Data Network, is considered.

C. METHODOLOGY

A literature review of the following subjects provided background understanding of internetworking:

- Broadband Integrated Services Digital Networks
- ATM
- Local Area Network technologies
- Fiber optic transmission systems
- Service provision in layered models
- The networking equipment industry
- The development, role, and state of networking standards
- Protocols for ATM internetworking
- Architectures of frame switches
- Frame and cell formats
- Throughput in contention-based networks

Analysis of frame and cell switching identified the strengths and drawbacks of each. An Intermodal Data Network (IDN) site visit provided insight into realistic management concerns. Interviews with the IDN Director and Senior Engineer clarified their goals in upgrading the IDN. The application of the ATM migration strategy to the IDN provided an opportunity to test the strategy through architectural analysis.

D. ORGANIZATION

Chapter II analyzes the uses and characteristics of frame switching. Frame switching network architectures are examined. The role of switching in conventional hub and router based networks is assessed. The various design approaches to frame switches are analyzed. Virtual LANs are assessed in terms of the strategic benefits they are purported to provide using multilayer frame switching.

Chapter III analyzes selected aspects of the Asynchronous Transfer Mode of data communications. An architectural analysis of the several roles ATM is capable of filling in the campus network is conducted.

The final chapter describes and analyzes the Intermodal Data Network, a complex multi-vendor, multi-protocol network. The use of a network modeling tool is explored. An analysis of the tradeoffs in implementing frame switching and cell switching in the Intermodal Data Network is undertaken. The ATM migration strategy is applied to the conventional IDN using an architectural analysis. A range of specific recommendations is made for upgrading the IDN to ATM through earlier implementation of frame switching.

II. ANALYSIS OF FRAME SWITCHING

A. CONTEXT

The explosive growth in desktop computing power is the primary driver of continually increasing network bandwidth requirements. Intel's establishment of a generational duration of 18 months has set in motion a blistering pace of development, deployment, use, and obsolescence. Furthermore, "Intel is by no means the sole instigator of the 18 month [cycle]. The type of industry -- competitive, unregulated, large number of competitors...dictates aggression. New products are the means to survival." (Budden2)

This rapid growth in hardware capability permits corresponding growth in application complexity, functionality, and bandwidth requirements. The rate of growth in bandwidth requirements is now exceeding that of processor performance improvements (SNCI).

The improved reliability of digital communications links and increased power available in PC-based servers and desktop machines is also fueling a trend toward distributed client-server systems. This increasingly popular paradigm is completely dependent on the interconnecting networks for its effectiveness. In fact, it can be said to be *Network Centric* (Budden1).

In the bridge and router based networks deployed in the 1980's, large Ethernet segments containing dozens of workstations were common. This architecture worked well at first, despite the complexity and performance impact introduced by the segmentation devices. During the early years of Local Area Network deployment when LAN-interconnect traffic consisted primarily of text-based information and most of the network traffic was destined for local addresses, this architecture consisting of relatively large segments worked well. The familiar 80/20 rule postulated that 80% of traffic would remain on the LAN, with only 20% forwarded to the enterprise backbone. That 80/20 split is undergoing an inversion. The second generation client/server architecture populates the backbone with distributed servers. With the edges of the network now

routinely requiring services across the backbone, we can expect to see 80% of the enterprise traffic on the backbone, with only 20% remaining within the local area.

The fundamental limitation of enterprise bandwidth lies in the fact that the transmission medium is shared by all users. A 10 Mbps medium may theoretically carry 10 Mbits of traffic, but any individual node will only see a fraction of this throughput. There is a wealth of research literature on this topic. (Stall94) provides an introductory treatment of the subject.

The classic approach to providing more bandwidth is to increase segmentation. This trend has been underway since the days of terminal-host communication when 1000 users might populate a segment. Recently LAN segment sizes have been dropping from dozens of nodes to only a handful.

The current approach to providing more bandwidth is to employ frame switching, using switching hubs. The primary benefit of using switching hubs is to increase bandwidth. A secondary benefit of switching hubs is that they facilitate the implementation of Virtual LANs.

Switching hubs operate at the MAC layer of the IEEE 802 Reference Model. They are protocol independent. They also employ one of two switching methodologies: either fast store-and-forward or cut-through. Both techniques yield switches with low latency, and thus high throughput. Design choices essentially make trade-offs between latency and throughput.

Comparing frame switching with routing from the OSI Reference Model perspective reveals some of the advantages switching offers. Routing occurs at the Network layer while switching occurs in the Data Link layer. The packets do not have to be examined at the Network layer for a routing decision to be made. Frame switching decisions are made based on the MAC frame itself.

Switches and routers provide some common functionality however. Both devices permit macro segmentation of the enterprise network, keep local traffic separate from

backbone traffic, permit high throughput, and facilitate end-user connectivity to the backbone.

The move from routing to switching also has topology and network management implications. Frame switching yields a flatter, simpler network, although it also requires large switching tables.

Higher speed shared media technologies also provide higher bandwidth on the LAN. Again, segmentation is used as a technique to ameliorate the drawback inherent in any shared media. As the bandwidth requirements increase, the new higher-speed technologies such as 100 Mbps Ethernet feel the same pressure to micro-segment the LAN. Switching permits us to further reduce the number of nodes per LAN segment.

ATM LANs are emerging into this context. ATM switching generates a price/performance ratio which is ten times greater than router based networks.(SNCI) ATM offers excellent scalability and the ability to implement virtual networking. An ATM network also promises to deterministic service required to support voice and video applications.

ATM is most appropriately being used in the campus backbone and vertical riser portions of private networks. The introduction of ATM into the campus backbone may reduce the life expectancy of FDDI as a backbone technology. Some early adopters have implemented ATM all the way to the desktop, but usually only at the 25 Mbps rate, and then only for certain high-end applications.

The potentially ubiquitous applicability of ATM to all pieces of the network (including the WAN) opens up additional possibilities. The current functionality provided by a LAN, keeping local traffic local and providing basic interconnection, may not be a necessary function in an all-ATM environment. The ability to create virtual channels, pipelines, and networks in a switching environment may relegate the LAN to the status of a familiar interface for network managers.

The development of ATM LANs is by no means complete. There are a multitude of unresolved issues involving network management, routing, addressing, route servers,

and LAN emulation, to name but a few. The IETF (Internet Engineering Task Force) and the ATM Forum are attempting to resolve and further develop these significant unknowns.

The virtual LAN feature promised by ATM is not an insignificant benefit. In fact, it may be the most desirable aspect of ATM, even when considering ATM's high bandwidth. The utility of the virtual LAN lies in the work force multiplier effect it has on the network management staff. LAN administrators spend the largest share of their budget on people. These people typically spend a majority of their time on adds, moves, and changes to the data network. Unlike the easier changes made to organizational voice networks, data network changes involve significant management effort. Any cost savings in this area are apt to be highly attractive.

ATM LAN switching serves to decouple the physical and logical implementations of the network. The hierarchical addressing scheme of the IP subnets and the MAC address of the workstation need not be reconfigured in the switched world every time an employee relocates their desk to another cubicle. Collocating the servers with the clients in the client-server architecture becomes unimportant. Careful design of physical segmentation is no longer needed to manage bandwidth. The Virtual LAN (VLAN) implementation is not without its price however. There is a large effort required to populate the VLAN (i.e., routing) tables. An automated, powerful management tool would be of enormous benefit to this effort. This is an area ripe for commercial development of management tools.

B. SWITCHING NETWORKS

The word *switching* is applied in three different contexts when discussing LANs. *Configuration Switching* provides the network manager the ability to use software to connect a port to another LAN.

> Some shared media hubs include functionality that allows the network manager to statically assign a repeater port to any of the LAN segments supported within the backplane of a hub. While the most descriptive term for this function is probably either *port assignment* or *port allocation*,

some vendors prefer the terms *port switching* or *configuration switching*. This static mode of altering configuration through a network management application should not be confused with LAN switching(SNCI1).

This relocation only *logically* moves the port to another hub. This capability is offered in such products as the Bay Networks Synoptics System 5000 (Lippis, Bay). This capability does nothing to improve performance however. The systems supporting configuration switching only facilitate a few local moves and changes. The difference between this configuration switching and VLAN reconfiguration is an issue of scalability.

LAN switching, or frame switching, does provide for performance improvement. The effect of frame switching is that of simultaneous bridging. The LAN frames are switched among the network nodes at native LAN speed. When building a new network without the requirement to incorporate existing infrastructure, the use of switches rather than hubs and bridges is preferable from the performance perspective.

ATM switching is characterized by the cellification of the LAN frames, prior to their being switched. LAN frames enter the ATM switch in their native format and are packaged into ATM cells before being forwarded to their destination. Once formed into cells, they are switched to their destination ports. ATM switching may occur at any point in a network. Workgroup switches, backbone edge devices, and wide area links may all use ATM technology.

1. Frame Switch Locations

There are various locations in the LAN where frame switching equipment may be placed. Switching hubs may be used to directly attach individual workstations. Each node will have a dedicated line connecting itself to the switch. Another option which supports LAN growth is to aggregate the nodes using conventional hubs, with each workstation having its own dedicated line into the conventional hub. These hubs may be stackable. The output side of the hubs are connected to a frame switch. While this greatly expands the number of workstations served by a single frame switch, it does introduce contention

within the shared-media hub. These switches may in turn be stackable, providing greater ability to keep local traffic off the backbone. Deploying LAN switches can be as simple as replacing hubs with switches.

The final use for frame switches within the LAN is as an edge device, located at the periphery of the backbone. The edge switch could be attached to an ATM backbone. It could thus provide LAN emulation software on the frame switching side, and would cellify the frames for transmission on the backbone side. A key concept in this analysis of LAN switching is that switched LANs provide a migratory path to ATM (Morency1).

2. Characteristics of Frame Switching

There are other features of switching hubs that make them attractive in a LAN environment. One reason for implementing switching hubs lies in their ability to readily connect a high density of power users. Another benefit is achieved when using the VLAN capability to create private LANs. A switch can also interconnect routers in a backbone. Finally, a switch can also interconnect a specific workgroup with a server, "front ending" the server with high-throughput connectivity.

The prices of switching hubs available today vary wildly, from $250 per port to $6,000 per port.(Lippis) Prices are declining on all these devices however. ATM switches are as much a part of the telephone industry as the data networking industry. This may translate to fast technology introduction and price competition. (Budden2)

Implementation of frame switching in today's client/server environment requires the network designers to understand their traffic flows. The traffic pattern can be generally typified as either client/server, or peer to peer. The switch may be connecting workgroup peers with a local server. This portion of the network may carry peer to peer as well as client to server traffic. Enterprise servers may be located outside the vicinity of the workgroup. The switch will then also carry client/server traffic through its backbone port.

Dynamic workgroup composition further adds to the dynamic nature of the traffic patterns in many organizations. When mission critical productivity in a client/server environment is directly related to the reliability and performance of the network, switching is a good fit.

3. Additional Applications for Switches

Switching hubs can be used to interconnect groups of servers. The servers would be equipped with unremarkable FDDI or Ethernet interfaces which could be attached directly to a switching hub. Management benefits may accrue from such a configuration. The servers can be centrally managed in one location, while still providing high performance. The configuration can be readily expanded by stacking more switching hubs, or by upgrading interface cards to higher speed technologies, such as FDDI or ATM.

Switching hubs can be used to augment the performance of a router-based backbone architecture. In this configuration the router retains the potential to become a bottleneck, interconnecting LAN segments with WAN connections and servers. The individual LAN segments realize performance improvement as switching hubs are added in a hierarchical fashion within the segments. However, reliance on the router maintains the status quo with regard to existing network shortcomings, e.g., lack of support for virtual LANs.

C. ARCHITECTURES OF SWITCHED NETWORKS

1. Distributed Backbone

The use of switches at the edge of the enterprise network form what we may refer to as a "distributed backbone" (Lippis). Populating the edges of the network with switches serves to "flatten" the backbone. The distributed backbone architecture is illustrated in Figure 1. Much of the traffic remains on the periphery of the network,

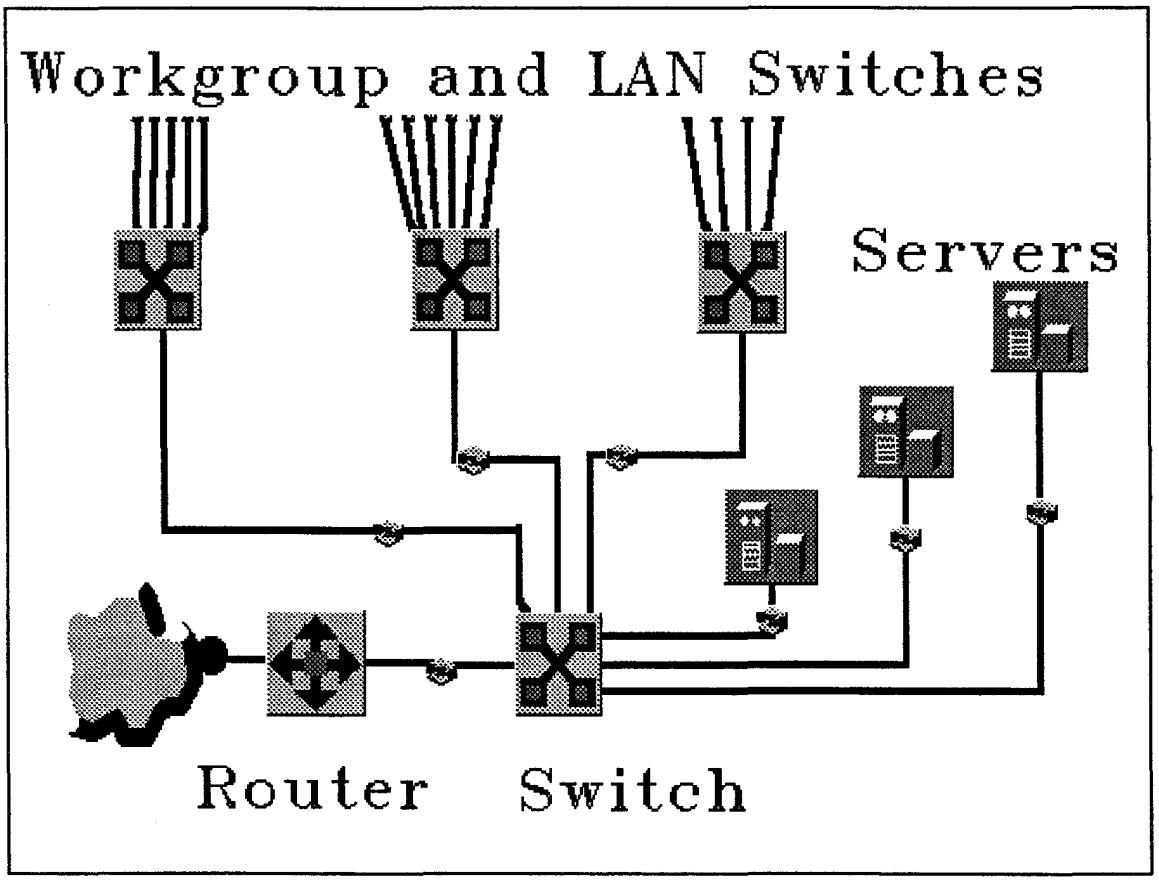

Figure 1. The distributed backbone architecture. (SNCI).

and does not necessitate transiting the backbone. One of the backbone switches will have a switched connection to a router for interconnecting WAN service. In this configuration the switching is moved close to the workgroup.

2. Distributed Backbone With Routing

Another distributed backbone architecture retains routing as the central mechanism for interconnecting LAN segments. The distributed backbone architecture relying on routing as the central mechanism for interconnecting LAN segments is illustrated in Figure 2. Switching is used to provide increased bandwidth in the segment and server locations.

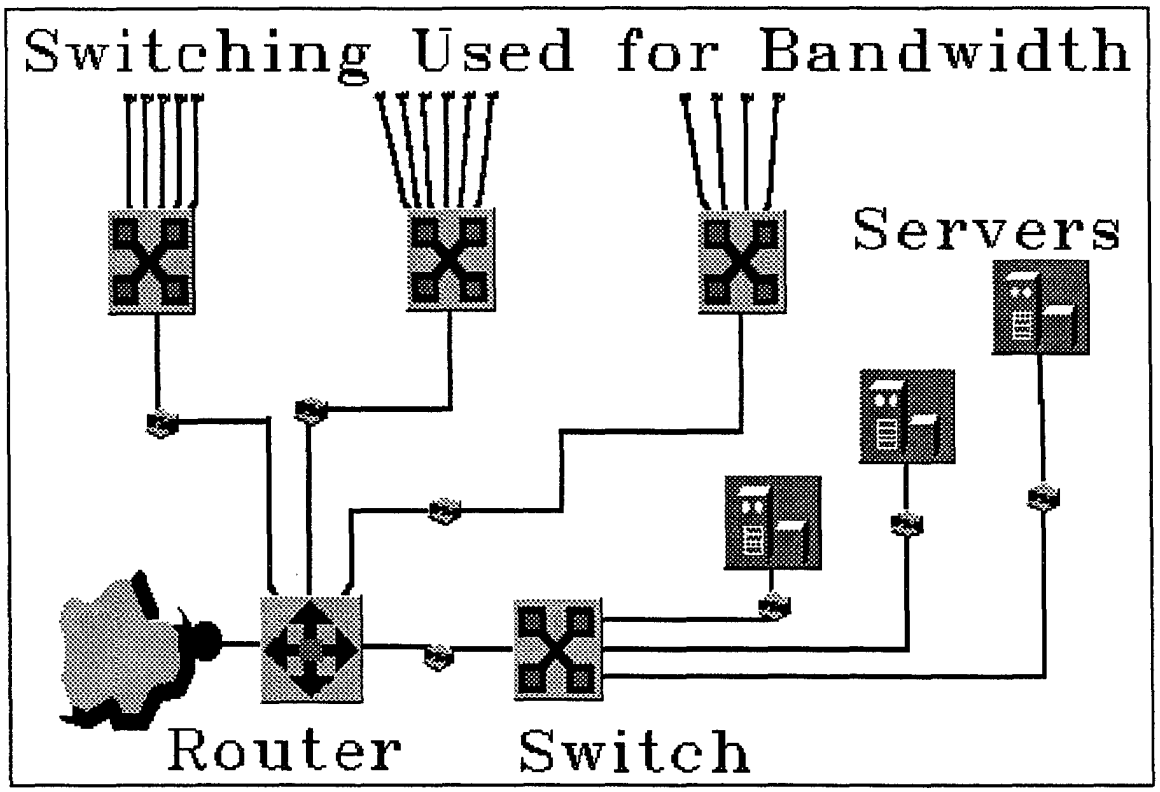

Figure 2. The distributed backbone architecture retaining routing as the central mechanism for interconnecting LAN segments.

The WAN connections can be made through the same router. Servers can likewise be directly connected to the router. This architecture is a likely candidate for an organization wishing to gain some familiarity with switching technology without abandoning its network investment. Switching hubs can be readily added to the edge of the network to alleviate congestion and provide increased bandwidth.

3. Distributed Backbone With Switching and Routing

A distributed architecture with integrated switching and routing is another, albeit more complex, option. In this integrated architecture, both routing and switching could be used on the network periphery. Figure 3 illustrates the distributed backbone architecture with distributed switching and routing.

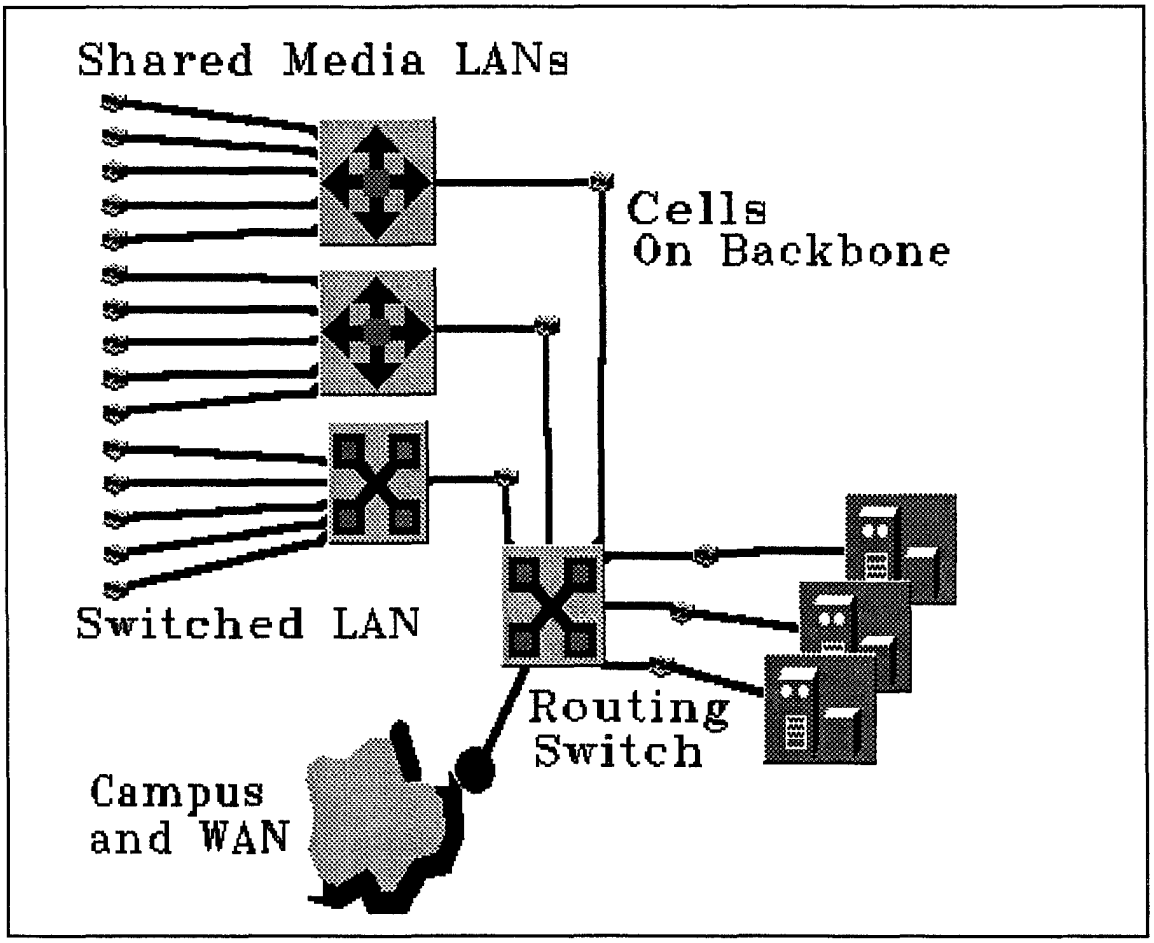

Figure 3. The distributed architecture with integrated switching and routing.

The backbone could consist of ATM switching. In this context we use the term "routing" as a logical function which may be distributed throughout the network, rather than as a function which only takes place within a "router" (Alles). The LANs at the periphery of the network can thus be combinations of switched LANs and shared media LANs. The distributed routers provide network layer services and concentrate traffic for the switched backbone.

Routing located at the backbone's edge permits easy interconnect of mixed LAN media (i.e., copper and fiber). Routing can also be used centrally to manage security concerns and to separate portions of the network from each other for control purposes.

4. Cross-Architectural Issues

Implicit in this architecture is the requirement for commonality of VLAN implementations across all manufacturers, protocols, and LAN technologies. Much work remains to develop such an open network. Although open solutions have yet to be defined, some manufacturers are making excellent progress with their proprietary approaches. Agile Networks, Inc. provide a switch which supports this architecture (Agile).

The Agile Networks ATMizer 125 Relational Switch combines frame switching and cell switching in a single device. The ATMizer 125 supports VLANs, which Agile Networks terms "relational LANs". The relational LAN management software automatically detects the end stations. The address of the LAN frames is examined to determine the membership of each virtual LAN. The use of an ATM switched backbone in this distributed backbone environment yields a "high capacity switching fabric that integrates existing endstations, routers, and wiring hubs into a single cohesive backbone network." (Agile)

This distributed backbone architecture provides three dividends. The first is increased bandwidth provided by an ATM switched backbone. The second dividend is the ability to take advantage of Virtual LANs, provided the manufacturers can standardize on an open solution. Finally, there is a significant reduction in personnel costs associated with being able to automatically detect and respond to adds, moves, and changes.(Lippis)

The increasingly large and scalable switching hubs on the market are leading to the emergence of a hubbed approach to an entire enterprise network.(SNCI) The high speed of the LAN switch's backplane ensures that stackables and modules can be added to support large numbers of end stations. The number of endstations involved can greatly exceed the numbers found in conventional concentrators and shared media hubs. This architecture greatly reduces complexity. It represents perhaps the least complex alternative available today for a high capacity network of a particular size. The advantage

of managing both physical and logical networks from one location is an added benefit. One disadvantage of this configuration is that cabling costs may increase; all workstations are home-run to the switching closet.

D. COST AND VALUE OF SWITCHING

The switching architectures outlined above offer significant value to an organization. "Switching provides better than a 10:1 price/performance ratio over shared media LANs and routers." (Lippis, SNCI) The use of switching permits taking advantage of the flexibility offered by virtual LANs. Reductions in the numbers of network operations personnel are possible due to the inherent ease of managing adds, moves, and changes. Alternatively, the operations staff can perform other functions which are slighted due to resource constraints. The massive movement from mainframe computing to desktop computing occurred when a price/performance improvement of 4:1 was achieved with Personal Computers (Lippis). If this ratio can be used as a guide for networking technologies, the same scale mass migration to switched networking may be upon us.

Two unresolved drawbacks remain unaddressed. The present status of enterprise networks as relatively open appears to be headed for the complicating world of proprietary solutions. The standards necessary for switched network implementors simply are not available yet. In this environment the demand for standards is outpacing the ability of standards bodies to create them. One can only wish that one's vendors are participating in the standards development process. This may be contractually enforceable.

The other issue involves cost of ownership. Switching really provides increased performance while only maintaining the high cost of network ownership. The increasing requirements for network traffic are fueled by client/server architectures and more powerful workstations. Network managers will have to continue their already high levels of resource commitment to implement switching networks which can keep up with the burgeoning demand.

E. ANALYSIS OF SWITCH ARCHITECTURES

Detractors claim that frame switching is just bridging. They are correct from one technical perspective, but frame switching also circumvents the problems inherent in bridging. Bridged networks are susceptible to broadcast storms from malfunctioning equipment. Networks can become flooded with useless packets. Isolation of these broadcast storms has traditionally been a challenge for managers of bridged networks.

The conventional approach to providing increased bandwidth availability in a shared medium network is to segment the bridge and router network. As the number of workstations per segment on a CSMA/CD (carrier sense multiple access/collision detect medium access control method) network decreases, the number of collisions drops. Probability of collisions decreases, and throughput rises dramatically when the number of workstations per segment drops below four. The throughput can be theoretically shown to rise from approximately 25 percent of medium capacity to 100 percent when the number of stations on the segment decreases from four to one (STALL94). When N active stations per segment approaches one, we may say that the network has been *microsegmented*. This microsegmentation effect is precisely the benefit achieved by LAN switches. LAN switches essentially microsegment to the limit of N=1.(SNCI)

LAN switches are thus a new class of networking product. They provide a dedicated network connection for each workstation or server, as if the LAN were microsegmented to the level of one workstation per segment.

LAN switches perform like a specialized high speed brouter (bridging router). The heart of the LAN switch is a very high speed backplane with large aggregate throughput. The internal architecture is characterized by multiple parallel paths. The LAN switch provides as many simultaneous connection paths as there are pairs of input/output ports in the configuration. The backplane rating should thus be at least a simple multiple of one-half the number of ports times the native speed of the attached media. Surprisingly, this is not always the case. Some first generation switches are limited in the backplane speed to aggregate throughputs which are less than a fully loaded switch would require. For

example, Xylan Corporation's Omniswitch featured a 640 Mbps backplane. If the switch were configured with 5 ASM-155FM ATM OC-3c Multimode Fiber Optic Switching Modules with two ports each, and each channel was saturated at capacity, the backplane capacity would quickly become overloaded. (Xylan)

LAN switches feature low latency through the switch. They also are beginning to have a low cost per port. Network Peripherals' EIFO is a cost-effective switching solution, offering FDDI and Ethernet capabilities for about $600 per segment. This FDDI price point is about half the price of other competitor's products. (Schireson)

Switching is topology independent. Network designers thus experience an additional degree of freedom with switching. Switches are not restricted to one specific LAN technology. In that regard they offer one of the benefits formerly only achieved through the use of a router. Switches can operate with any broadcast or token-passing LAN, and can support a heterogeneous diversity of LANs.

There are four major switch architectures in use today: cross bar, shared memory, shared bus, and cell bus.(SNCI)

1. Crossbar Switches

Crossbar switches were the most common first generation implementations of LAN switches. Kalpana was an earlier implementor of LAN switches, relying on the crossbar architecture. Other vendors similarly used the crossbar architecture for their first generation switches.

The crossbar switch architecture uses a switching matrix to switch *N* input ports to *Y* output ports. The switching matrix is controlled by a central controller. The frames entering the switch have their address read, and a switching path is established immediately. The LAN frame is actually departing the switch before it is completely received into the switching matrix. Figure 4 illustrates the conceptual basis for the many-to-many port mapping of a crossbar switch architecture.

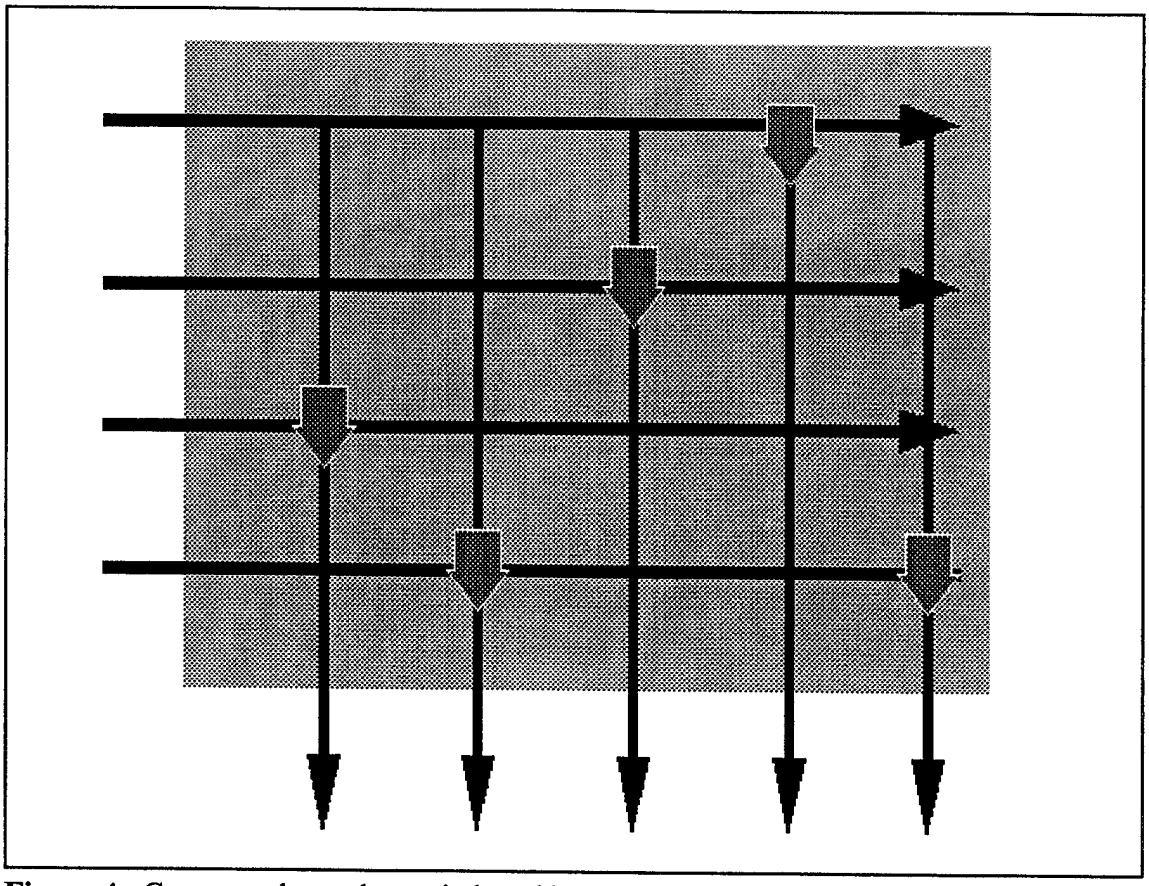

Figure 4. Conceptual crossbar switch architecture. (Partridge).

The performance of crossbar switches and other frame switches has been tested recently by Scott Bradner of Harvard University. The 1995 Switched Ethernet Evaluation results may be obtained at (SNCI1).

The crossbar switch has several desirable attributes. The architecture is satisfactory for switches with a small number of ports. The switches exhibit low latency characteristics. The crossbar architecture is a good fit with a peer to peer, high density workgroup interconnection environment.

The crossbar architecture's shortcoming lies in its poor scalability. Although Digital Equipment Corporation's GIGASwitch/FDDI system has been scaled up to 55 ports, the configuration is very expensive. (Lippis, DEC) The crossbar architectures are a poor choice when ports are operating at different speeds. They do not support multi-speed switching without a large amount of buffering. The inclusion of the required

buffering essentially makes the switch operate like a shared memory architecture switch. The crossbar switch is thus best suited to a like-to-like network architecture.

Crossbar switches are widely available from a variety of manufacturers today. Examples include (SNCI, Lippis):

- Cisco Systems' (Kalpana) EtherSwitch
- Digital's GIGASwitch
- Bay Networks' (i.e., SynOptics) switching engine
- Xylan Corporation's OmniSwitch and PizzaSwitch (Xylan2)

2. Shared Memory Switch Architecture

As its name implies, the shared memory switch architecture is characterized by the presence of a pool of high speed shared memory. Packets enter the switch in the normal, serial fashion. They then enter a serial to parallel multiplexer which performs two functions. The packet headers are copied and sent to the control module while the entire packet proceeds into the shared memory area. The control unit determines which output port the packet should be switched to. It then directs the shared memory to pass the packet to the parallel to serial demultiplexer for serialization and switching to the appropriate output port. The parallel to serial demultiplexer then transmits the packet to the proper output port.

The shared memory switch also has several desirable attributes along with some drawbacks. This architecture can yield very high throughputs due to the combination of the high parallelism with the lack of delay inherent in memory to memory packet transfer. Maintaining this high transfer rate is dependent on designing a sufficiently large bandwidth between the input port and the memory, and the memory and the output port. Manufacturers such as 3Com claim "full-wire-speed forwarding on all ports for all configurations - no exceptions(LANplex)".

The architecture can support multi-speed interconnect requirements. In addition to 10 Mbps to 10 Mbps connections, both 10 Mbps to 100 Mbps and 100 Mbps to 100 Mbps connections are feasible. This architecture can be implemented at a relatively low cost by building the specific shared memory management logic into hardware. This is achieved through the use of ASIC (Application-Specific Integrated Circuit) technology, a critical core competency area for successful switch manufacturers.

The primary drawback to the shared memory switch architecture is that the bandwidth between memory and the input/output ports can become a bottleneck. Available implementations display an upper limit of approximately 16 ports per switch. There is also some industry interest in hybrid designs combining the attributes of shared memory and shared bus architectures. Shared memory switches are also widely available from a variety of manufacturers today. Examples include products from (Lippis):

- 3Com Corporation (LANplex 2500 series)
- Bay Networks (28000)
- OST, and
- Cabletron

3. Shared Bus Switch Architecture

The shared bus switch architecture features a passive backplane which switches ports through a Time Division Multiplexing (TDM) scheme. The backplane itself is a shared medium. The attached ports essentially broadcast their traffic across the backplane.

The shared bus architecture offers many strong arguments for its design. The media independence of the backplane provides this design with excellent flexibility. The backward compatibility of this design enables the network manager to leverage the installed base of existing equipment.

The architecture has a strong combination of high bandwidth and configuration flexibility. Highly segmented LANs and VLANs are a good fit with a shared bus approach. This architecture is an excellent choice for use as a collapsed backbone device, and is most commonly used for that purpose.(Lippis) Bay Networks' SynOptics System 5000 is a switch which meets this demand. The shared bus switch architecture will also support the use of ATM as the backplane media.

The fundamental design limitation of the shared bus approach is the bandwidth of the backplane. As discussed earlier, the aggregate bandwidth of the incoming ports should not exceed the backplane's capacity.

Shared bus architecture switches are available from a number of manufacturers today. Examples include(SNCI, Lippis):

- Chipcom's Online
- Network Systems/Bytex 7760
- Digital's DECSwitch 900
- Cabletron MMAC Plus
- SynOptics LattisNet System 5000
- Optical Data System's Infinity
- 3Com's LANplex 6000

4. Cell Bus Switch Architecture

The cell bus switch architecture connects the individual logical LANs by means of a *cell switching bus*. By using cellification on the LAN switch backplane, the cell bus architecture is ready for upgrading to ATM technology. This architecture is a good choice for switches which will be used as edge devices. The consistent form of the cells being switched affords the design considerable flexibility.

Many of the cell bus switches have been implemented with an additional, parallel bus known as the *management bus*. The management bus approach is not unique to the cell bus architecture however. The implementation of a separate dedicated management bus ensures that management functions are available during periods of high switch utilization or congestion. This addresses a significant irony in network management; devices are often difficult to reach and manage precisely when that capability is needed the most.

Cell bus switches are also widely available from a variety of manufacturers today. Examples include products from (Lippis):

- Centillion
- Xedia
- FORE Systems
- LANNET, and
- ONET

F. VIRTUAL LANs: THE STRATEGIC DRIVER

Migrating to switched networking will eliminate sharing the interconnecting media. There is another compelling argument for the use of switches: switched networks support *virtual workgroups*. The structuring of these flexible workgroups into *Virtual LANs* is a fundamental, and highly significant benefit of switched networking.

The management of adds, moves, and changes is the largest component of the people cost associated with operating a network. The people portion is likewise the most expensive piece of network ownership. The ability to implement virtual LANs directly addresses network management's largest cost sector. This sector has been estimated to represent 23% of the people costs, and 10% of total network ownership costs (SNCI).

From one perspective, virtual LANs operate as if they were broadcast groups. Groups are composed of the workstations which share some common work task, organization, or interest. The formation of virtual LANs based on some business-oriented reason is a departure from the conventional practice. A workstation is traditionally added to a LAN based on the physical location of the computer. LANs are organized geographically by building, department, floor, cubicle and the like. These configurations are driven by the practicalities of physical LAN implementation, and to provide convenience to the LAN administrators. These are not business-driven reasons.

Unfortunately, current implementations of virtual LANs are achieved through proprietary methods. These proprietary schemes typically rely on MAC addresses for subnetwork definition. Detractors would rightfully argue that this drawback is of overwhelming significance. Non-proprietary open solutions are in the process of standardization however. It is left to the network manager to ascertain which of the many vendors is offering a solution which appears to be most in compliance with the emerging standards.

New "second generation" VLAN implementations promise dynamic autoconfiguration of subnetworks (Lippis). This capability would help reduce the operational expenses associated with adds, moves, and changes.

Virtual LANs can be implemented in switching hubs and ATM switches. A robust network management capability is thus developed as a by-product of the implementation of VLANs.

VLANs decouple the physical network address from the logical address. This decoupling permits the moves and changes to occur with regard for the rules and limitations imposed by the IP addressing requirements. The address administration takes place within the devices, and can occur automatically through autodiscovery routines.

Each workstation may thus belong to a number of VLANs. This is an excellent match with the way business is actually conducted. An end user may be part of the local

workgroup VLAN, a larger departmental VLAN, a wide area interest group VLAN, and their own additional personal-interest VLANs.

The VLANs also support collaborative work flows, one of the more widely regarded information technology developments. These technologies permit groups of people to share work in progress. VLANs also permit the logical collocation of the components of client/server computing. All servers could be logically collocated on one segment of a VLAN to provide the network administrator with simplicity and clarity. The physical location of the servers then becomes moot and server farms with concomitant broadband access may not be needed.

The VLANs may then be extended across the enterprise network and the wide area network. These extensions will require routers and ATM switches.

A key concept of the migration strategy advocated by this analysis is that *the routing functions can be achieved through the use of frame switches which contain ISO Layer Three functionality.*

The support for VLANs in ATM switches is only just emerging in mid-1995. The network manager must again rely on their own analysis of the fit between the emerging standards and the vendor's offerings. Network managers must attempt to avoid proprietary solutions which result in dead-end technologies. For this reason many industry observers are cautioning network managers to buy homogeneous solutions from vendors with a large market share.

III. ANALYSIS OF ATM CELL SWITCHING

A. ATM IN CONTEXT

Information technology developments are rightfully criticized for promising to provide a technical cure for many existing problems. The difference between the hyperbole and the reality frequently results in a shortcoming of the form "promised too much, too soon." Actual experience might indicate that technical solutions frequently substitute one (known) set of problems for a new and less familiar set. Strategic Networks Consulting, Inc. surveyed network managers prior to the publication of the first ATM buyer's guide. The survey respondents clearly indicated that "future-proofing" their networks was their second highest priority, only surpassed by their desire to increase bandwidth. The purported abilities of ATM to support improvements in network management, improve network security, and establish virtual workgroups were met with skepticism by the respondents. (Morency)

Much has been said and written about ATM. Its promises are being viewed with great skepticism by many practitioners. *Another Technology Mañana* and *Assimilate, Takeover, and Merge* are two of the acronym's translations suggested by detractors. ATM is nonetheless being embraced by the networking community. In a survey conducted at Networld+Interop in March 1995 hundreds of networking professionals were asked what backbone and desktop connectivity they were presently using, and what connectivity they would be using in two years. Thirty five percent indicated that their organizations anticipated using 155 Mbps ATM in their enterprise backbones, and eight percent predicted they would have 155 Mbps ATM to the desktop. The results are presented in Figure 5.

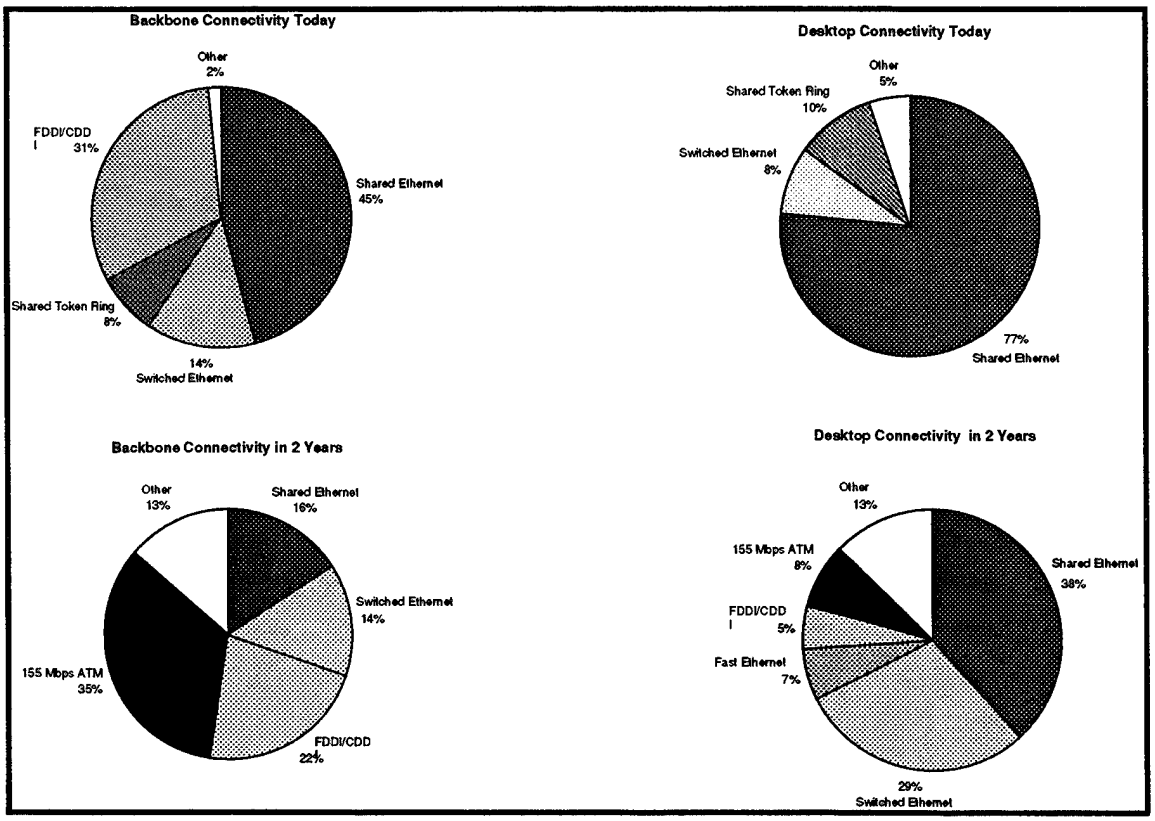

Figure 5. Backbone and desktop connectivity today and in two years. (Survey).

ATM is still in the early adopter stage of deployment. Practicing network professionals must not become enamored with the way networking "should be", but must focus on real opportunities and real problems. This is particularly relevant to any contemplated migration to ATM.

B. VIRTUAL CHANNELS AND VIRTUAL PATHS

1. ATM Clocking

The term *asynchronous* in ATM refers to the indeterminate nature of the "ownership" of the next arriving cell. This is quite apart from the notion of guaranteed rates of service, and is not in contradiction of that attribute of ATM. An ATM link can carry numerous logical connections simultaneously. The next cell of a given logical connection will arrive at some time in the future. The use of the term asynchronous is thus analogous to its historical meaning; the next cell arrives on its own schedule.

Cells of one form or another are always arriving however. "The preceding cell predicts nothing about the next cell except when it will start -- immediately after the current cell ends." (Flanagan) TDM mechanisms carrying ATM cells always operate synchronously. They derive timing explicitly from a master clock or recover timing implicitly from the network itself. ATM networks commonly use fiber optic transmission media. This fiber is a serial, TDM link.

2. Aggregation of Connections

Routing in ATM networks uses two hierarchical levels of identifiers to describe the *channels* and *paths* which cells will take through the network. These are defined in ITU-T Recommendation I.113, *Vocabulary of Terms for Broadband Aspects of ISDN.* A single link between two devices in an ATM network is called an ATM Peer-to-Peer (APP) connection. An APP will typically carry several *virtual paths* and *several virtual channels*. A single physical circuit may contain several virtual paths which in turn contain several virtual channels. This bundling permits intermediate switching nodes to switch the entire path or channel without examining every cell within that aggregation.

The *Virtual Path Identifier* (VPI) and *Virtual Channel Identifier*(VCI) comprise the first three bytes of the five byte ATM cell header. These identifiers are "used to describe unidirectional transport of ATM cells associated by a common unique identifier value." (Siemens) The VPI and VCI are simply logical identifiers for switch input port to output port mapping (SNCI). The VPI consists of 8 bits, and the VCI uses 16 bits.

ATM is a connection-oriented transmission method. All cells in a specific connection are guaranteed to take the same route to their destination. Cells thus can not overtake each other, but discarded errored cells may be missing upon reassembly at the destination. Each link in any specific connection is assigned an identifier which remains unchanged for the duration of the connection (Siemens).

> Virtual connections are built up by linking APPs end-to-end. The way these linkages are defined, set up, and taken down determines how connections and paths are implemented. Thus a call setup process is

needed to assign addresses, and relate each address at each node to the desired logical connection. An ATM node uses addresses to direct a cell to the APP which is the next segment of the connection. (Flanagan)

3. Permanent and Switched Virtual Connections

Addresses that are assigned and maintained for long periods of time create *Permanent Virtual Circuits* (PVCs). Circuits set up on the basis of a real-time request are known as *Switched Virtual Connections* (SVCs). Manual implementation of PVCs may take network operations personnel hours to set up or modify. This delay may not be problematic if service is expected to be provided for a long period of time. The PVC is somewhat analogous to a leased line. Implementation of the dynamically negotiated SVCs is lagging PVC rollout.

C. ATM EQUIPMENT

1. Characteristics of ATM Switches

Abstractly, ATM switches are comprised of four primary components. The first is the set of interface modules. These components make the connections to the switch from possibly several different sources. These ports can serve legacy frame-based technologies such as FDDI, Token Ring, Fast Ethernet, or Ethernet. Alternatively, they can be cell-based interfaces. These would typically support connections at optical carrier speeds, and comply with the User Network Interface specifications.

The second component is the processing module. This element performs several cell processing functions. Cellification, scheduling, and connection management are carried out in the processing module. A benefit of cells is that they work very well in parallel processing devices. Variable length packets are thus purposefully eliminated early in the processing module's sequence of operations. This ensures that remaining processing will take advantage of hardware's strength in processing homogeneous cells very rapidly. The module also contains the switching fabric itself.

The third element of the ATM switch is its high speed trunking. Previously-switched cells are now aggregated at high speeds. Typical trunk speeds are in the T-1 to OC-12 range. Several trunks may be served by a single switch.

The final component of an ATM switch is the device management of the switch itself. This function spans the other components beginning with cellification and continuing through processing to the exiting trunked data.

ATM switches may be characterized into three major segments: small workgroup switches, enterprise (or campus) switches, and network edge switches . Most of the present activity involving ATM in the user's environment is occurring in the LAN. John Morency has estimated that as of March 1995 there were approximately two thousand ATM switches installed. These switches typically were configured with two to four ports per switch. (Morency)

Small workgroup switches might operate at the 25 Mbps rate and serve a small group of power users working with an application such as CAD. Edge switches serve a function somewhat similar to a gateway, accepting inputs in legacy frame formats, breaking open the frames and performing cellification, and finally injecting the cells into the cell-based network. The term *enterprise* in *enterprise switch* is a bit misleading. The enterprise switch serves a single campus setting such as a high-rise building, rather than providing all service to a single small organization with WAN connectivity requirements.

Both the workgroup switches and the enterprise switches are configured with interfaces supporting most media types. They include:

- multimode fiber at 100 and 140 Mbps
- single mode fiber at the SONET 622 and 155 Mbps rates
- coaxial cables at DS-1 and DS-3 rates, and
- twisted pairs at up to 155 Mbps (over distances such as 20 feet).

2. Characteristics of ATM Adapters

Applications which are to be transported using ATM must have their information placed into an acceptable ATM format prior to transmission. Workstations which are to be connected to ATM networks meet the cellification requirement by using network interface cards or adapters. The ATM Forum is developing standards for ATM APIs which are presently adapter-specific. A primary problem with ATM adapters is that their costs have remained high as their production numbers remain low. Strategic Network Consultants reported that S-bus adapters are now selling for one seventh of their price two years ago. There is no reason that ATM adapters for client workstations will not become commodity items like conventional NICs. (SNCI)

The adapter is comprised of three main portions: the interface with the workstation bus, the ATM segmentation and reassembly (SAR) engine, and the interface with the physical layer of the ATM network. The ATM engine performs the SAR to correspond to the requirements of the User Network Interface (UNI) and the ATM Adaptation Layer (AAL) in use.

> Server engines are a different animal. Inherently, they require significantly more sessions of simultaneous Segmentation and Reassembly. The large number of logical connections greatly increases overall complexity and requires large amounts of high speed RAM. (Flanagan)

The ATM Forum is contemplating adding virtual LAN emulation functionality directly on the adapter.

The selection of adapters to be placed in the workstations warrants careful consideration. Users of equipment from more than one vendor are typically restricted to using PVCs. There are some exceptions where switch manufacturer's SVC implementations are interoperable with those of other vendors. There is an ATM consortium at the University of New Hampshire which focuses on interoperability testing. The group expects greatly improved adapter-to- adapter SVC compatibility to become available during 1995. Additional information is available on the Internet. (UNH)

Adapters also display two approaches to the segmentation and reassembly function. Early ATM adapter designs performed this function in hardware while most new switches rely on software. (SNCI)

D. ANALYSIS OF ATM SWITCH ATTRIBUTES

1. Edge Switches

ATM edge switches provide the normal array of cell switching functionality but perform an additional function -- conversion from legacy frame formats to ATM cells. The port interfaces solely support frames, unlike the larger switches which also support entering cells. Backplane speeds are in the 200 Mbps to 2.5 Gbps range. The prices of edge switches has recently declined with lower end models in the $8,000 to $12,000 range, and others extending to approximately $35,000. (Morency, SNCI)

Edge switches typically support ATM network interface trunk speeds of OC-1 (58 Mbps), OC-3, and TAXI.

2. Workgroup Switches

Workgroup ATM switches usually have an upper limit of approximately 16 interface ports and reside in the $12,000 to $25,000 price range. The interface modules typically support the TAXI, OC-1, OC-3, OC-12, ATM-25, T-1, and T-3 rates (SNCI, product literature). All switches support PVCs (Permanent Virtual Circuits). They also support various proprietary interfaces. There is an increasing mass of support for the UNI 3.x standards. The 100 Mbps TAXI rate, which uses the FDDI physical layer, was intended to be the lowest ATM speed. TAXI is diminishing in importance while the ATM-25 and T-1 rates increasingly appear to be where the future of the workgroup switches lies. There is a perception that ATM-25 was promoted by IBM to justify their continued support for their 16 Mbps Token Ring.

Workgroup switch processing modules use backplane capacities which are typically less than 2.5 Gbps. Trunking configurations include support for OC-3, T-1, and T-3 speeds. The trunking ports also support PVCs, the Network to Network Interface

(NNI), and the ATM Forum's Private Network to Network Interface specifications. The switches do not incorporate any high availability features.

The workgroup switch is generally viewed as almost a niche solution for CAD or medical imaging users. However, industry observers believe large amounts of revenue could be generated by bringing ATM to the desktop for two applications: video conferencing, and World Wide Web browsers. The price for ATM components would have to drop significantly for ATM to become ubiquitous. The prices are currently plummeting, having already dropped 6-7 times in the past two years. (Morency) The pricing of ATM-25 is beginning to be competitive with 100 Mbps shared media systems. The difference between these two advertised rates might make Ethernet seem more attractive, until the decreased throughput of a contention-based MAC method is considered. ATM can provide guaranteed data rates from end to end. Ethernet will never be able to match that guarantee.

3. Enterprise Switches

Enterprise (or campus) switches share many of the same attributes as the workgroup switches, but offer improved processing power and a wider variety of trunking options. The prices are also larger. Campus switches cost from $25,000 to $300,000, depending on configurations. Backplane speeds are in the 2.5 Gbps to 10 Gbps range.

The interface modules supporting TAXI and OC-1 rates are becoming less common. OC-3, OC-12, ATM-25, T-1, and T-3 rates are usually available (SNCI, product literature). OC-3 is emerging as the lowest interface speed routinely offered. All switches support Permanent Virtual Circuits and various proprietary interfaces. They also provide support for the UNI 3.x standards. The campus switches introduce the much more reasonable OC-12 trunk speed and usually retain the OC-3 and T-1 and T-3 rates.

E. SWITCH SELECTION CRITERIA

There are several important criteria to consider when selecting an ATM switch. The general classification of the switch (outlined above) must be determined. This decision is dependent on the number of ports required, and the estimated total capacity

desired in the switch. The expansion capabilities of the switch should be given a relatively heavy weighting. ATM is clearly an evolving technology and the bandwidth requirements that current applications require may be rapidly supplanted by much larger rates.

Another important consideration in switch selection is the availability and compatibility of the ATM adapter to be placed in the workstations. While industry offerings are available for the common bus implementations, the immature standards issue immediately becomes a concern. Many of the features a network manager may be attracted to could be proprietary, and not supported by other ATM components from other manufacturers. The ATM market is not yet amenable to a robust mix-and-match approach to network design. An additional example of the compatibility concern is evidenced in the lack of interoperability among proprietary PVCs and SVCs. In general, there are differing levels of compatibility seen in switch-to-switch connections versus switch-to-adapter connections.

Public sector ATM users should consider placing the burden of proof of compatibility on the vendor or system integrator. One way to accomplish this is through the requirements written into an *Indefinite Delivery, Indefinite Quantity* (IDIQ) contract. Live demonstrations of interoperability can be mandated for products the vendor proposes to sell to the government. Interoperability with systems outside the scope of the contract are more problematic. The range of interoperability demanded of the vendor, and the methods used to test that interoperability, must be both well-designed and achievable.

Compatibility with the existing installed network is a key element of any realistic approach to ATM. There are probably very few networks which will make wholesale upgrades to ATM. Managers of those "turnkey" transitions would benefit greatly from the opportunity to design a homogeneous, fully compatible solution without regard to interoperability with existing infrastructure. The implementation of virtual LANs illustrates the problem. LAN emulation using the guidelines contained in RFCs 1483 and 1577 works well in IP implementations but has experienced problems when using IPX

(Morency). The ATM Forum's Version 1.0 of the ATM LAN Emulation specification will now permit vendors of internetworking and adapter cards to offer standardized LAN emulation products.

The network management capabilities of the ATM switches is also a major inconsistency from vendor to vendor. The MIBs supported by the vendors vary. This is a key concern; the promise of interoperable ATM to ease the network management burden is one of its chief attractions.

The ability of a large vendor to provide a total ATM solution is an attractive characteristic. This single-source approach can potentially eliminate the network manager's incompatibility problems. The primary drawbacks of such an approach include the perceived shortfalls in the vendors' proprietary network management functions and the inability of the network manager to substitute alternate equipment with more attractive features.

The non-standardized features are not simply address and device management issues. The manner in which congestion is avoided, the numbers of Virtual Paths and Virtual Circuits supported, and the quality of service features supported are of concern to the network manager. The ATM Forum offers additional criteria for use to consider when selecting ATM equipment.

There are additional substitution problems in the public sector. The process of substituting new components for those contained in the government's contract does not keep pace with technology changes. The contracts become progressively obsolete. One solution is for the government to award contracts more frequently. This increases the logistical effort required to support the larger variety of hardware and software however.

F. CONNECTION MANAGEMENT

Connection management refers to a signaling protocol used to set up and tear down PVCs and SVCs. The protocol provides the mechanism by which traffic can be re-routed and accounting information can be captured. Other properties of the connection

may be negotiated such as delay, delay variability, error rate, and the guaranteed throughput capacity (Flanagan).

Implementing PVCs in commercially available products can be straightforward. An example is Fore Systems' ForeView management product. Network operators can use a graphical user interface to set up unique VCI and VPI values for each switch between the source and destination. This is nonetheless a labor- intensive accounting exercise and helps to explain why mesh topologies using many switches are uncommon; they would be a network management nightmare. (Morency)

Connection management can be implemented in either a centralized or distributed fashion. In a centralized approach the connection management is implemented in software positioned with one of the ATM switches on the ATM backbone. This approach is used by Bay Networks, Inc. and in the LattisCell, Newbridge, and Hughes switches (Bay, SNCI).

A distributed connection management approach places the connection management responsibility throughout the ATM cloud. This architecture is implemented in FORE Systems' ForeThought software distributed in each ASX-200 switch. The software provides what FORE describes as Bandwidth Management. Service levels are guaranteed on a per virtual circuit basis. Strategic Networks Consulting, Inc. has evaluated the advantages and disadvantages of the centralized and decentralized connection management schemes. Their analysis is presented in Table 1. (FORE, Morency)

G. MULTILAYER SWITCHING CHALLENGES

ATM poses a long list of implementation difficulties. Most of these problems have several proposed solutions. The widespread adoption of standardized solutions is of critical importance in order for ATM to realize its promises. Partridge writes, "Today's internetworking protocols have proved quite capable of transmitting at gigabit speeds, with some modest tuning." However, he continues, "Today's protocols are not perfectly suited to supporting multimedia applications. These applications require performance

Evaluation	Centralized	Distributed
Scale	May become bottleneck	Intelligent switches, router based model
Low complexity	Single source resource allocation	Software distributed to every switch
Ease of software upgrade	Single source	Every switch needs upgrading
Connection setup	Setup time may suffer	Faster: resources allocated locally
Best disaster recovery	Single point of failure	One switch cannot bring down net
Cost	Less	More

Table 1. Distributed vs. Centralized connection management. Advantages are highlighted. (SNCI).

guarantees that today's protocols generally cannot provide: guarantees about maximum delay or minimum bandwidths." (Partridge)

LAN emulation over an ATM WAN illustrates another problem. ATM LAN frame encapsulation is outlined in RFC 1483. That guidance addresses how to envelope various Packet Data Units (PDU) within cells e.g., 802.5 frames (1483). Similarly, RFC 1577 addresses the operation of IP and address resolution within a Logical IP Subnetwork (1577). However, vendors implement these services in differing ways. Proprietary extensions are common.

Although the ATM Forum has more work remaining, progress is encouraging. Version 1.0 of the ATM LAN Emulation specification developed by its LAN Emulation Sub Working Group was released on May 8, 1995. This key standard opens the door to one of the most important end-to-end ATM capabilities.

> ATM LAN Emulation was designed to allow existing networked applications and network protocols to run over ATM backbone networks in a standardized manner. It supports not only end-systems directly-attached to ATM networks but also those connected via Layer-2 bridging devices. ATM LAN Emulation is important because it allows multiple LANs to co-exist on the same physically interconnected ATM network. (Forum)

If a similar degree of standardization evolves in the network management area, the performance/cost benefits promised by ATM may begin to emerge. Internetworking with ATM is indeed highly complex.

IV. APPLICATION OF THE MIGRATION STRATEGY

The use of frame switching as a stepping stone to full ATM networking is a migration strategy which provides increased bandwidth now while mitigating the risk inherent in a wholesale adoption of ATM. Frames switches of the *edge switch* variety discussed earlier provide ATM interfaces as well as legacy frame format interfaces. These frame switches provide greatly increased bandwidth now and ensure that the network manager is well positioned to introduce ATM into the enterprise network without scrapping the existing infrastructure.

This chapter examines a complex multi-vendor multi-protocol network, the Department of Transportation's Intermodal Data Network. The Intermodal Data Network interconnects the many heterogeneous networks existing throughout the Department of Transportation.

An analysis of the tradeoffs in implementing frame switching and cell switching in the Intermodal Data Network (IDN) is undertaken. Frame switching and cell switching architectures are then applied to the existing conventional IDN internetwork architecture. Finally, recommendations are made for migrating to switching networks.

A. DOT NETWORKING ENVIRONMENT

1. The United States Department of Transportation

The United States Department of Transportation (DOT) is an association of federal regulatory agencies with its headquarters in Washington D.C. With over 107,000 employees and a budget of $36.5 billion in fiscal 1993, the DOT has responsibility for nine administrations whose jurisdictions include highways, civil aviation, mass transit, railroads, commercial use of outer space, the merchant marine and the safety of waterways, ports, highways and oil and gas pipelines. The DOT also oversees the United States Coast Guard, a branch of the armed forces responsible for search and rescue at sea, marine environmental protection, and the enforcement of maritime laws and treaties. (DOT)

For purposes of this analysis, we may consider that the department consists of 12 separate components. Some of these organizations are known as *Operating Agencies* (OAs), and others directly support the Secretary of Transportation. From an internetworking perspective, we may regard each of these entities as a separate autonomous organization, each with its own network. We will refer to them as OAs. They are the:

- Federal Aviation Administration (FAA)
- United States Coast Guard (USCG)
- Maritime Administration (MARAD)
- Federal Highway Administration (FHWA)
- National Highway Traffic Safety Administration (NHTSA)
- Federal Railroad Administration (FRA)
- Federal Transit Administration (FTA)
- Saint Lawrence Seaway Development Corporation (SLSDC)
- Bureau of Transportation Statistics (BTS)
- Research and Special Programs Administration (RSPA)
- Office of the Inspector General (OIG)

2. Decentralized Control

The DOT Operating Agencies design, build, and manage their own networks. There is no centrally managed effort to standardize the protocols, hardware, software, media, functionality, or other aspects of their networks. The OAs have pursued their own procurement strategies on their own schedules. The result is a truly heterogeneous

multi-vendor multi-protocol network. Each agency manages and operates their own network in a manner of their own choosing.

3. Heterogeneous Networks

The DOT's administrative networks in Washington, D.C. reflect the wide variety of possible implementations. DOT agencies with an operational mission, i.e., the FAA and the Coast Guard, additionally operate separate multi-billion dollar mission critical networked systems.

The terminals and servers include a wide variety of devices including PCs, Apple computers, Sun SPARCstations, and Silicon Graphics workstations. The diversity of applications using the DOT IDN is illustrated by the variety of electronic mail systems in use. They include ccMail, Qmail, MSMail, BTOS Mail, and Groupwise E-Mail.

Various network operating systems are also in use. Many of the networks rely on Novell NetWare, but others use Banyan Vines, Appletalk, and even an OSI implementation used in the proprietary architecture of the Coast Guard's Unisys BTOS system.

Equipment from a significant number of vendors is also in use. Hardware such as routers, uninterruptable power supplies, concentrators, remote access servers, sniffers, modem pools, and hubs bear the labels of mainstream companies such as Cabletron, Cisco, Fibronics, Network General, 3Com, Wellfleet, Synoptics, APC, and US Robotics. WAN connectivity is equally diverse. The FAA maintains an AT&T-provided frame relay connection from a router on the main FDDI ring to a router and Ethernet LAN in Oklahoma City. The Volpe Transportation Systems Center in Cambridge, Massachusetts is connected to the IDN via an AT&T-provided T-1 line. The primary Internet access is provided commercially by Advanced Networking Systems through a 56 Kbps Bell Atlantic-provided connection.

B. NETWORK DIAGRAMS

The symbology shown in Figure 6 will be used throughout the IDN network diagrams in this chapter. The symbols were generated using (CACI).

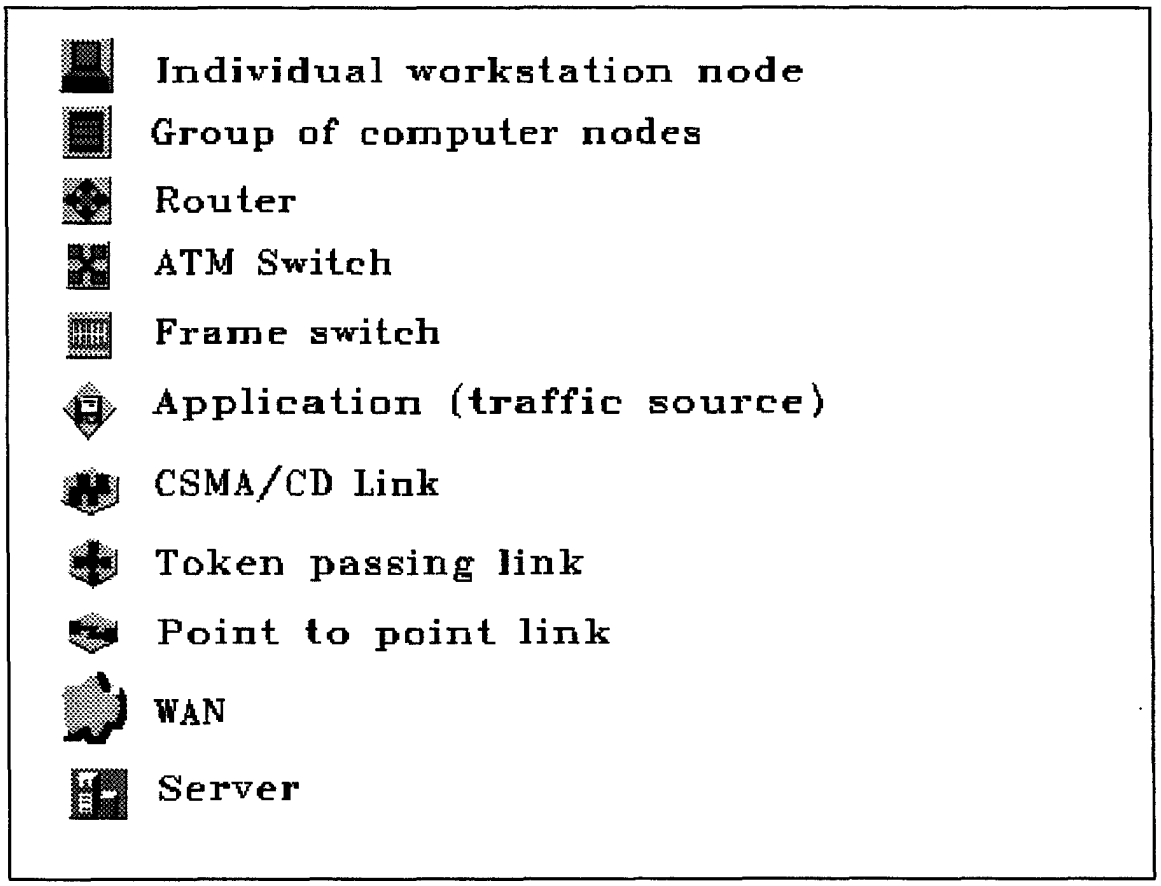

Figure 6. Key to symbols used in IDN diagrams.

C. IDN FDDI MAIN RING CONFIGURATION

The main ring of the IDN backbone is located within the Nassif Building in Washington, D.C., primarily within the building's penthouse. The backbone is a Dual Attached Station (DAS) FDDI ring. A DAS ring actually consists of two counter-rotating fiber rings connected to each device. The FDDI ring is populated with approximately 20 routers. With the exception of two Cisco products, the routers are 3Com Corporation's NETBuilder II model. Figure 7 provides a simplified diagram of the network emphasizing the main ring populated by routers, one representative subnetwork in the lower right corner, the Internet connection in the upper right corner, and the World Wide Web,

Gopher, and Domain Name servers at the upper left. Note the ability to keep Internet traffic destined to and from the servers from transiting the main ring.

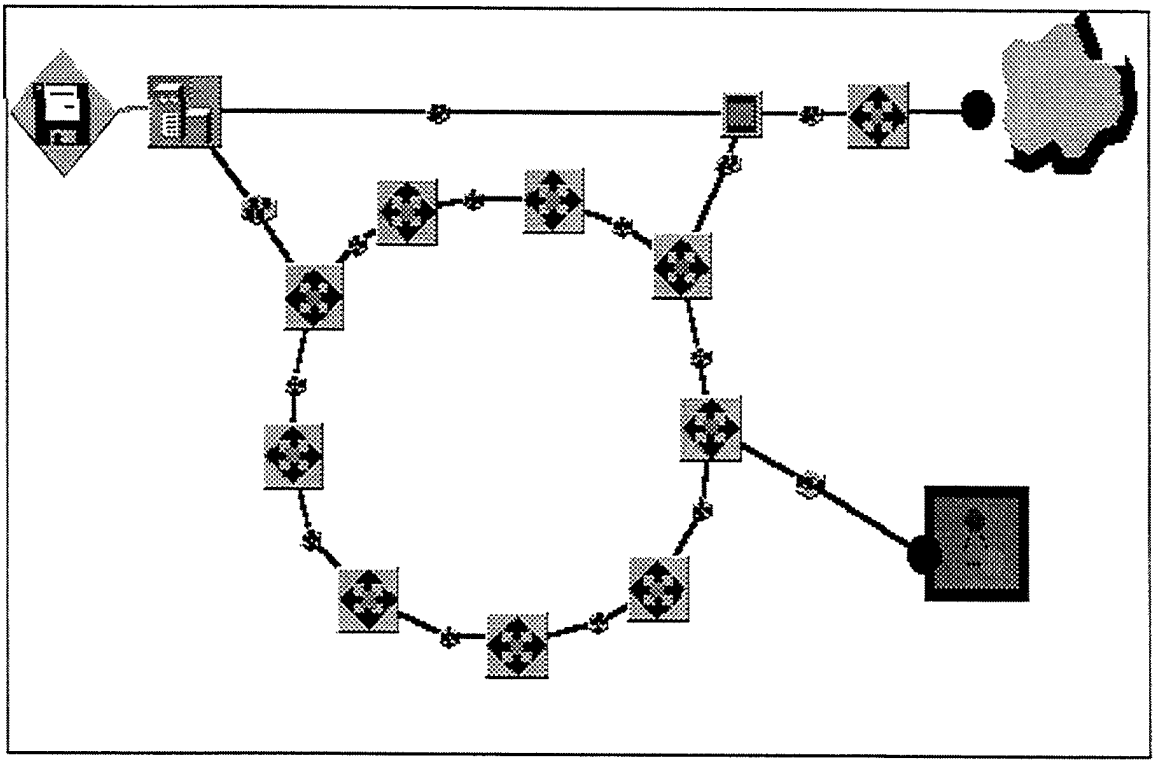

Figure 7. Simplified diagram of the IDN FDDI main ring with Internet WAN connection, one subnetwork, and servers. Internet traffic destined for the servers need not transit the ring.

The entire backbone network is managed with Cabletron's *Spectrum* network management software. Spectrum uses a database application, SpectroServer, and its GUI-based application SpectroGraph. The IDN operates the Spectrum system on a Silicon Graphics Indigo workstation. The primary SpectroServer and SpectroGraph, named *Savanna,* is located in the network management center. Four additional SpectroGraphs are dispersed throughout the OA networks. The networks with SNMP-addressable SpectroGraphs belong to the FHWA, FTA, MARAD, and RSPA. Additionally, a full SpectroGraph/SpectroServer suite is located on the FRA's network.

D. CHARACTERISTICS OF THE NETBUILDER II

The 3Com Corporation's NETBuilder II is a high performance multi-protocol router. The device combines the functionality of both a translational bridge and a router. It features a modular design to permit flexible configuration alternatives to meet the specific needs of each piece of the network. The modular design also provides an upgrade path: components can be replaced as the technology progresses. The router supports ATM, Ethernet, Token Ring, FDDI, and a wide variety of WAN connections. 3Com claims that the basic architecture offers a smooth growth path to ATM and Fast Ethernet. The 800 Mbps backplane speed indicates that this claim may be valid.

The NETBuilder II uses custom ASICs and RISC processing. The incorporation of multiprocessor modules permits additional RISC processing power to be added to any critical interfaces. The WAN connections support an extensive array of services including channelized T-1 or E-1, Switched 56, and ISDN. The router provides robust support for IBM SNA protocols, integrating SNA support in a multiprotocol environment. It performs this function by encapsulating SNA data in IP packets in accordance with the Data Link Switching standards of RFC 1434. The NETBuilder II supports Advanced Peer-to-Peer Networking (APPN). The router also supports the major LAN and WAN protocols such as TCP/IP, IPX, XNS, OSI, DECnet Phases IV and V, Banyan VINES, Appletalk, and others. (NBII)

E. MODELING IDN TRAFFIC

An extensive effort was undertaken to model traffic flows on the IDN using the CACI Products Company's *COMNET III* network modeling software, release 1.1n. This section describes the author's model, and the subsequent section details the results of the simulation. A diagram of the actual model is presented in Figure 8.

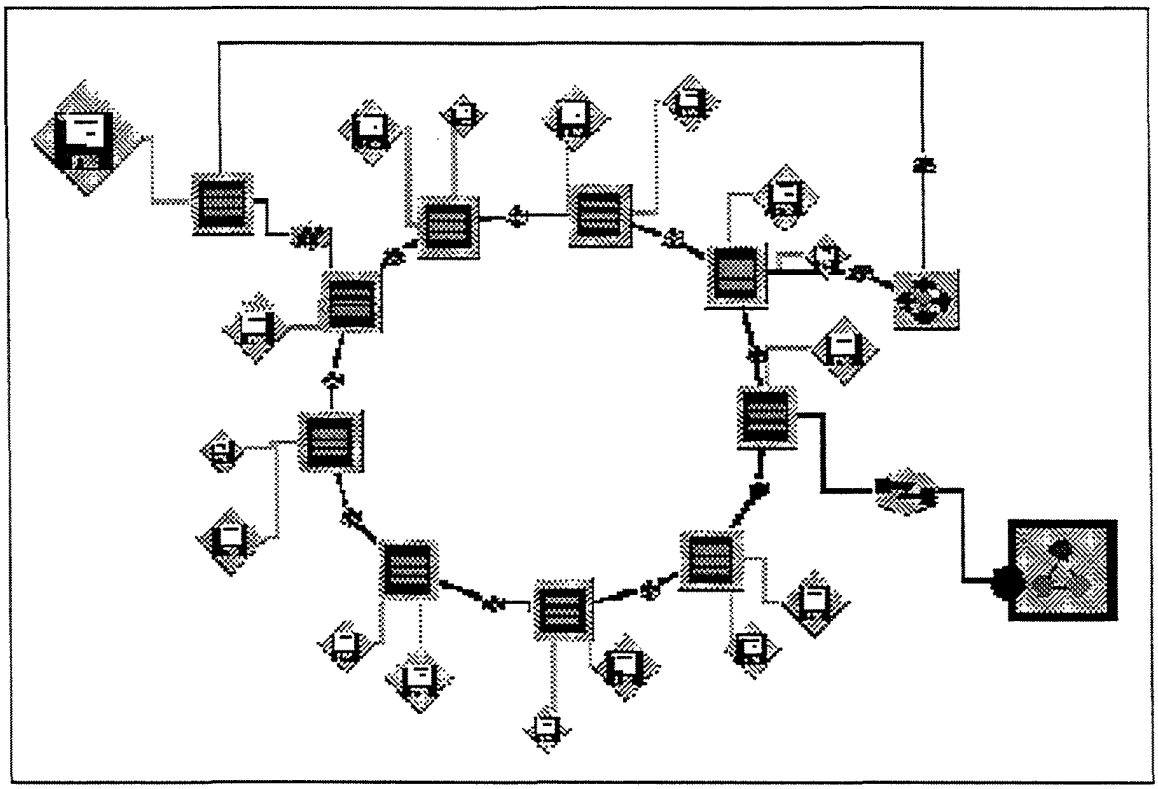

Figure 8. IDN FDDI Main Ring Model. Traffic-generating applications are adjacent to the group of computer nodes/router nodes. The modeled subnet is in the lower right. The CSMA/CD links connect the servers and the subnet to the main ring.

The IDN FDDI Main Ring Model was developed as a simplified image of the actual network. The main ring consists of nine routers in a Dual Attached Station FDDI ring. Three of the NETBuilder II routers play a key part in the model. These routers and their modeled characteristics are listed in Table 2. The remaining six routers in the model also generate backbone traffic.

Traffic for these six remaining subnets was aggregated at each router node. The subnetwork's traffic is generated by the applications directly linked to the router and is represented as an aggregated bit rate received and forwarded at the router's level on the main ring. The router named NBII-R3-2 actually has a subnet attached to it in this model.

Traffic from this subnet is routed packet-by-packet to the main FDDI backbone and then to its final destination. The subnet is named *Typical Subnet* and is depicted in Figure 9.

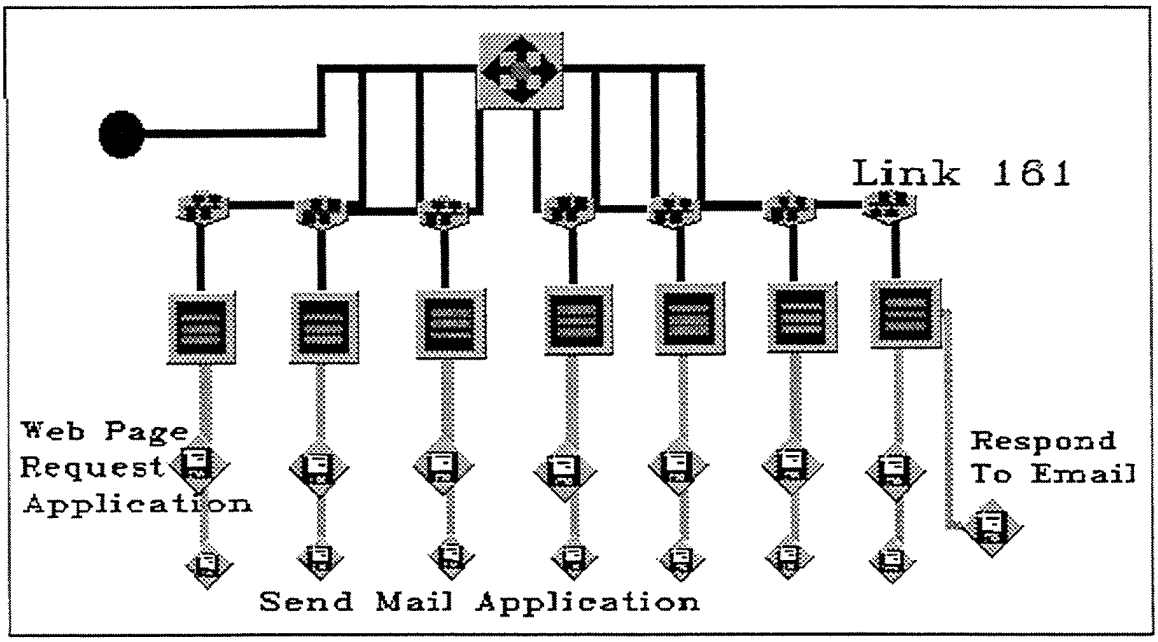

Figure 9. Configuration of the *Typical Subnet.* The subnet uses a contention MAC and generates Email and World Wide Web page requests. The node on Link 161 also responds to Email it receives.

F. MODELING RESULTS

The modeling effort was not entirely satisfactory. A rigorous statistical analysis identifying the locations of the most heavily congested network segments was not developed. However, the information which was obtained corresponds to the architectural analysis.

The author believed that a simulation of several hours of network operation would provide sufficient data for subsequent analysis. The model was run on Intel 486-based PCs with 16 Mbytes of RAM. Approximately the first 23 minutes of network operation were rapidly simulated in a few minutes. However, the computer required approximately

Router / IP Address	Primary Role	Interfaces	Customer
NBII-R1-2 152.120.100.34	Provides FDDI ring access to primary and secondary DNS, World Wide Web, Gopher, and I Mail Servers	port 3	DOT M70 802.3 LAN
		port 4	WAN to Volpe
		port 5	Network Management Center 802.3 LAN
		port 6	802.3 LAN
		port 7	DOT M90, M50 802.3 LAN
NBII - (PL 300) / NBII-SW4-2-P8 152.120.110.201	Makes connection from IDN to Internet	port 3	FAA HDN
		port 4	USCG/DC
		port 7	DSU - frame relay WAN to Oklahoma City
		port 8	Internet connection via 10BaseT
NBII-R3-2 152.120.100.38	Representative connection for heterogeneous LANs	port 3	FAA 802.3 LAN
		port 4	OST 802.3 LAN
		port 5	FTA 802.3 LAN
		port 6	RSPA 802.3 LAN
		port 7	BTS Vines network
		port 8	FAA 802.3 LAN

Table 2. Significant router and port configurations modeled in IDN FDDI Main Ring Model.

thirteen hours to simulate the final minute of operation. The disk drive access light appeared continuously lit during this time. This may indicate that the RAM was fully used, and subsequent steps in the simulation required continuous hard disk drive reads and writes. The author halted the simulation after 1,572.7 seconds (~26.2 minutes). A more satisfactory simulation is outside the scope of this thesis.

The short duration of the simulation constrained the attempt to generate heavy loading on links of interest. The effects of heavy traffic with resultant collisions would have indicated the location of bottlenecks within the network. The limited simulation lengths prevented large numbers of collisions from developing. In the 26+ minutes of simulation only three CSMA/CD collisions occurred, effecting six frames.

G. LINK DELAYS AND UTILIZATION

The COMNET III software can be configured to generate a report specifying the link delays and link utilization for selected links in the model. These links carry connectionless traffic. Link loading data was generated for several links of interest. Data was also obtained for several other links in the FDDI Main Ring Model, particularly in the vicinity of the servers. The loading on these links, in terms of percentages of their 100 Mbps capacity, was in the very low single digits. These results reveal the limited complexity of the model. Data was extracted for three CSMA/CD links. These links are described in Table 3.

Link Designation	Link Location
Link 108	Connects Web, Domain Name Servers, and Gopher Servers to Main Ring
Link 159	Connects Main Ring to Typical Subnet
Typical Subnet Link 161	Connects Single Workstation to Subnet's Router

Table 3. Modeled links carrying connectionless circuit traffic for which link loading data was extracted.

Not unexpectedly, the simulation clearly showed that the heaviest loading occurred on the 10 Mbps CSMA/CD links connecting the servers to the Main Ring, and the subnet

to the Main Ring. All three collision episodes occurred on Link 108. Six frames were involved in the collisions. Each episode was resolved on the second attempt.

The extracted data reflecting the loading on Links 108, 159, and 161 is displayed in Table 4. The identification of the link is presented along with the number of frames each link delivered. This is the number of frames removed from the output buffer of the transmitting node and delivered to the input buffer of the receiving node. The transmission delay is calculated as the difference from the time the frame is created at the input node to the time it is delivered at the end of that link. The modeling software can produce this data both in cases where the frame size exceeds the packet size, and when the packet size can envelop the entire frame. The inclusion of a standard deviation value should not be interpreted to imply that the delay times follow a normal distribution. The data is tabulated as an average value, a standard deviation, and a maximum value for each link.

		Transmission Delay (ms)		
Link	**Frames Delivered**	**Average**	**Standard Deviation**	**Maximum**
Link 108	1972	0.489	0.546	1.508
Typical Subnet Link 161	5297	0.522	0.561	1.221
Link 159 FROM Node	2162	2.003	3.226	7.813
Link 159 FROM Typical Subnet	5793	7.014	2.145	7.813

Table 4. Transmission delays and loading data generated by the simulation.

H. TRADEOFFS IN IMPLEMENTING SWITCHING IN THE IDN

The principal tradeoff in implementing frame and cell switching in a conventional network such as the IDN is fundamentally one of cost versus capability. The large

increase in LAN bandwidth achievable through the migration from a shared media LAN to a switched LAN incurs a large increase in network ownership cost.

The price of a high quality Ethernet hub is significantly lower than the price of an Ethernet frame switch. For example, 3Com's LinkBuilder FMS II Series of hubs are stackable and can be configured for any Ethernet media. The hubs are fully compatible with similar units designed for Token Ring and units designed for remote access. The cost of a 12 port twisted pair media hub, including the SNMP management module, is $ 1,674. (Industrial)

The same number of ports can be obtained in the Ethernet frame switch manufactured by ONET. The ONET LB2000 is a stackable 12 port twisted pair Ethernet frame switch with fixed latency and compatibility with other high speed transmission methods. Configured with SNMP management capability, the switch costs $ 10,161. (Industrial)

This comparison indicates that migrating to a frame switching network will require a capital investment which may approach six times that required for a conventional hub-based LAN. However, the bandwidth improvement is much more than six-fold. The microsegmentation effect of switching networks completely eliminates the shared media problem. Every node receives 100% of the media speed.

A node on an Ethernet LAN with 20 users would be fortunate to approach 1 Mbps throughput. Data throughput rates below 200 Kbps would be likely, based simply on dividing the available bandwidth among users and disregarding contention effects.

Frame switched networks have even more bandwidth readily available to them. Many of the switches offer the ability to configure pairs of ports for full duplex rates. A 10 Mbps Ethernet switch could be configured for 10 Mbps to and from a server. The same is true for 100 Mbps Fast Ethernet implementations. This full duplex 100 Mbps option is an attractive alternative for connecting servers to the network.

Finally, many frame switches are upgradable to ATM. The interfaces on the network side of the switches are already available in ATM configurations. This option

permits the sequential migration from frame switching within the LAN to cell switching on the backbone. This arrangement can be followed by increasing uses of cell switching. The next migratory step is the incorporation of cell switching on the LAN side of the switch. These ATM modules are also available today, but are not available for every switch from every vendor. These modules would be most appropriate for switches which already use cell switching in their backplanes.

The ability to replace an existing frame switching module with a cell switching module is a tremendous benefit to upgraders. The simple addition of an ATM expansion module is likewise possible. An ATM expansion module added to an existing frame switch provides the ability to connect the server directly to the frame switch using an ATM link. The switch aggregates and cellifies the incoming frame switched traffic and maps the cells to the cell switched output port.

I. SWITCHING ARCHITECTURES APPLIED TO THE IDN

There are two portions of the DOT networks which are strong candidates for the introduction of frame and cell switching architectures. The first is the centrally managed Domain Name, World Wide Web, and Gopher server segment. The second is within the subnetworks themselves. Specific recommendations for upgrading these portions are included in Section J below. A recommendation to upgrade the IDN FDDI Main Ring is notably absent.

1. Applying Switching To Server Connections

The IDN server "farm" is the primary candidate for migration to frame and cell switching. The applicable servers and their functions are:

- Mailstorm - the Secondary Internet Domain Name Server
- Depository - the World Wide Web and Gopher Server
- Kiosk - the Primary Internet Domain Name Server
- Compass - the External Internet Domain Name Server

These servers are interconnected through several critical hubs. The interconnecting links are 10 Mbps Ethernet connections. The area of the IDN formed by these hubs, servers, and interconnecting links holds several possibilities for migrating to switches. This area is depicted in Figure 10. Hub 4 carries traffic for Mailstorm, Depository and Kiosk. This traffic originates from outside the DOT network. The low speed of the commercially provided Internet connection, 56 Kbps, makes an ATM switch upgrade an unnecessary option. However, as useful information is added to the DOT servers, traffic from the Internet can be expected to increase. The Internet connection could be upgraded to a megabit rate. A switched network connection would then begin to make sense if the link were heavily loaded.

Hub 4 itself could be replaced with an ATM switch or a frame switch right now. All three servers could be connected to the switch with Ethernet, Fast Ethernet, or ATM links. The servers would then be able to generate responses without experiencing frame collisions on their presently shared media. Hub 5 serves a similar function for users located inside the DOT network. This hub is directly providing connections to Kiosk and Depository. This architectural design is an important part of the IDN firewall configuration. For this reason integrating more of the connections into a single switch should not be considered. Hub 5 could readily be replaced with an ATM or frame switch. The network side of the switch could be directly connected to the NBII-R1-2 router with full duplex 100 Mbps connections. This would be an excellent option to visibly improve service to the DOT's own IDN "customers".

Hub 6 is perhaps an even better location for a switch because it also serves the network engineering support staff (not depicted in Figure 10). The location of a new technology here would raise its visibility to the network support staff. This would serve to ensure that close monitoring occurred while familiarity with the new technology develops.

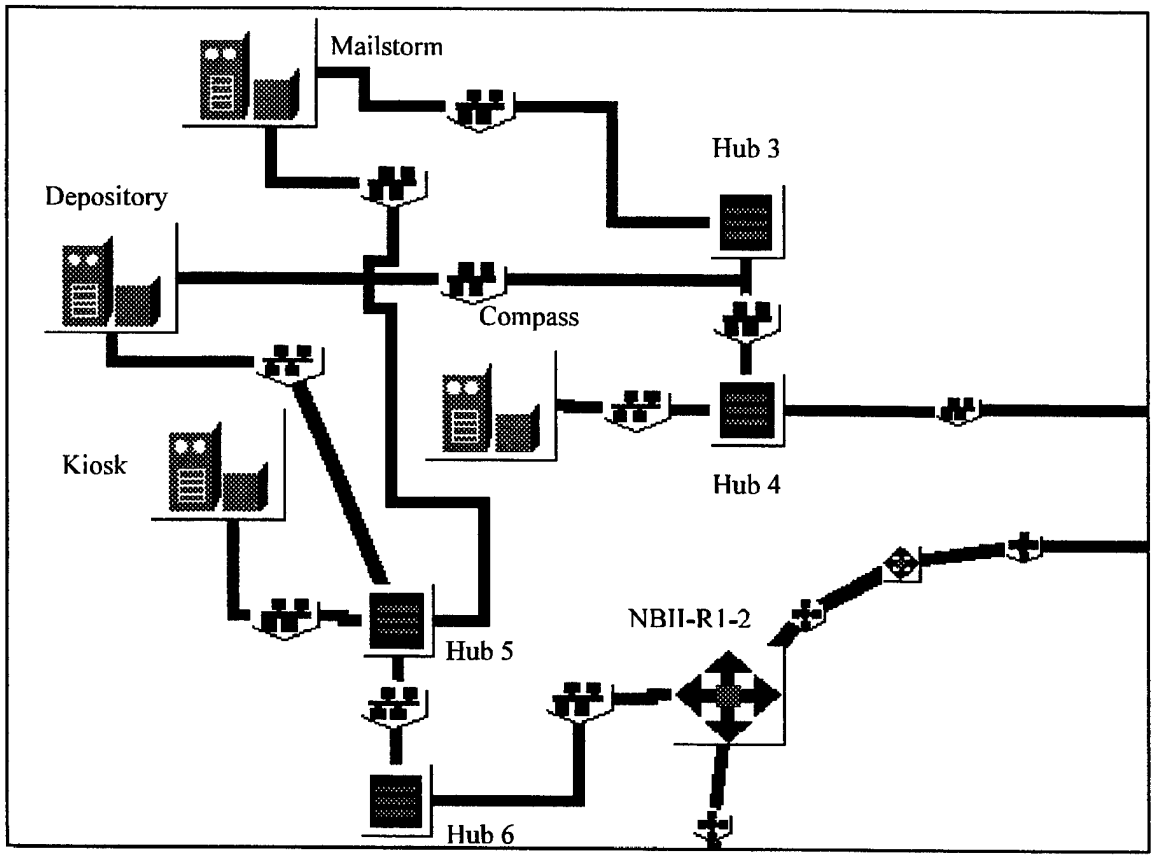

Figure 10. Detail of the IDN server topology. Contention techniques are used for medium access control of all links displayed outside the main FDDI ring.

Despite these architectural advantages, Hubs 5 and 6 present particular real-world dangers to the IDN support staff. All DOT IDN traffic destined for the servers passes through these devices. Any failures or shortcomings during the early stages of implementing the new technology would be highly visible to the DOT IDN customers. Any problems would serve to lessen the credibility of the network support staff. The resultant intra-organizational tensions would require a lengthy period of time to fade away.

A highly attractive alternative is to place a switch on the LAN used by the network support staff itself. This option would permit Hubs 6 and 5 to continue to reliably carry the server load. The network support staff could gain familiarity with switching technology without risking the high availability that IDN users have come to expect.

An option with more associated risk is to replace router NBII-R1-2, located on the IDN FDDI Main Ring, with an ATM switch. FDDI interfaces are readily available, and each server could be individually connected to the ATM switch. Ring-based traffic would pass through the switch as before, but the possibility of congestion occurring on the Ethernet server links would be eliminated.

2. Applying Switching To Subnetwork LANs

The use of frame and cell switches within the subnet LANs is a more straightforward decision. The incorporation of ATM switches within the subnets is most likely not warranted at this time, but may be fielded if desired. Frame switches are a timely choice however. These switches can be used in place of any existing hub. They should be targeted first for replacement of those hubs at the highest levels of cascaded hub architectures. LAN switches in these locations will thus be switching traffic which is already concentrated at lower levels.

J. RECOMMENDATIONS FOR MIGRATING TO SWITCHING NETWORKS

Frame switching should be applied to the IDN in the LAN backbone areas of the subnetworks. The link delays and utilization data contained seen in the modeled Typical Subnet indicate that this architectural approach is sound. While some LAN traffic does transit the IDN backbone, a significant portion of the traffic generated on the Typical Subnet remains local to the Operating Agency's LAN. Frame switches interconnecting an OA's own Ethernet, FDDI, and Token Ring LAN segments will thus bring an immediate bandwidth improvement to local-based applications traffic.

The IDN support staff should purchase a pair of ATM switches and install them on their own subnetwork. This will provide the network engineers with significant exposure to cell switching technology prior to implementing it IDN-wide. A chief procurement strategy in choosing these ATM switches should be to focus solely on the standards which are supported by the ATM switch. Proprietary enhancements should be disregarded

during the equipment selection process. This approach is particularly important in the IDN's multi-vendor multi-protocol environment.

Upgradable frame switching should also be applied to the IDN-managed servers. World Wide Web page requests can be expected to sharply increase when an initiative to place official copies of Requests For Proposals (RFPs) on the Web server is completed. These RFPs are substantially more interesting to the business community than the "brochureware" which the Web server currently offers. This application is indicative of the greater utilization the network will experience in the near term future.

These upgradable frame switches should be very carefully selected with ATM in mind. ATM modules should be gradually incorporated into these switches as demand rises and ATM links are implemented. This phased migration to ATM will mitigate the risks as the ATM standards evolve, and familiarity with the technology grows.

The recommended architectural solution for subsequent upgrades to the server connections is to connect an ATM switch to all the externally-viewable servers. Router NBII-R1-2 should be considered for the first replacement of a Main Ring router with an ATM switch. Finally, the IDN should procure an ATM WAN connection only when that service is both readily available and warranted.

These recommended upgrades of the IDN to frame switching and ATM provide other key benefits. ATM is growth-oriented and flexible. ATM standards ensure that the IDN will be increasingly interoperable with other government agencies and private sector organizations. The strong interoperability attributes of ATM also apply to developing technologies. Adopters of ATM will not be left with obsolete networks before they achieve their return on investment. The combination of frame switching and ATM switching will help the IDN control the uncertain risks of a wholesale upgrade to ATM.

K. CONCLUSION

The demand for bandwidth is growing at a rate faster than workstation processing power (Lippis). Networking professionals are faced with the challenge of providing large

increases in available bandwidth. One aspect that favors success in this endeavor is the decreasing incremental costs of bandwidth.

> The predominance of fixed costs in transmission link systems leads to declining cost per unit of information moved ($/bps) as the volume of information increases up to the maximum capacity of the link. Also, the low incremental cost of constructing higher-capacity means that they generally have lower unit costs of information movement than lower-capacity systems. For example, a system with a data rate capacity of 1.544 Gbps [Mbps] will generally have a lower cost per bps than a transmission link system with a 64-kbps data rate capacity. This is because the basic construction, installation, and maintenance costs of the two systems using any single medium is approximately the same, even though the larger system has roughly 24 times the capacity of the smaller system. (Keen)

The enterprise or campus backbone is the portion of a general purpose business network most likely to require ATM. There is no clear requirement for ATM speeds in the IDN backbone today. The FDDI Main Ring is currently only experiencing *six percent* loading. (Saliba) However, the use of ATM switches in the backbone would certainly place the IDN in a strong position for future requirements growth.

Further ATM standards are expected in the next 18 months. The era of demand-driven ATM deployment in enterprise networks may be commencing. Unprecedented price and performance points may provide organizations with networks which contribute toward meeting business objectives in new and significant ways. "It is clear that ATM will play a central role...." (Alles1)

LIST OF REFERENCES

[1483] Heinanen, Juha. 1993. Multiprotocol Encapsulation over ATM Adaptation Layer 5. Network Working Group, Request for Comments 1483. URL:http://ds.internic.net-/rfc/rfc1483.txt. July.

[1577] Laubach, M. 1994. Classical IP and ARP over ATM. Network Working Group, Request for Comments 1577. URL:http://ds.internic.net/rfc/rfc1577.txt. January.

[Agile] Agile Networks, Inc. 1995. ATMizer Relational Switch. World Wide Web. URL: http://www.agile.com/product-info/atmizer.html.

[Alles1] Alles, Anthony. 1995. *ATM Internetworking*. San Jose, California. Cisco Systems, Inc.

[Alles] Alles, Anthony, William Hawe, Nicholas J. Lippis III, and Ron Toth.. 1995. *Will ATM Switches Replace Legacy Routers ?* Conference Discussion . Networld + Interop 95, Las Vegas, Nevada. Cassette 150327-C26. La Crescenta, California: Audio Archives International.

[Bay] Bay Networks, Inc. 1995. Product Information. World Wide Web. URL:http://-www.baynetworks.com/Products.

[Budden1] Buddenberg, Rex. 1995. Network Centric. Lecture notes.. World Wide Web. URL:http://dubhe.cc.nps.navy.mil/~budden/lecture.notes/network-centric.html.

[Budden2] Buddenberg, Rex. 1995. Unpublished notes to the author.

[Budden] Buddenberg, Rex. 1995. Computer Networking and C3I Systems for Emergency Services. World Wide Web. URL:http://dubhe.cc.nps.navy.mil/~budden/book/table_contents.html.

[CACI] CACI Products Company. 1995. COMNET III User's Manual: Planning for Network Managers. La Jolla, California: CACI Products Company.

[DEC] Digital Equipment Corporation. 1995. Switching Product Family. U.S.A.: Digital Equipment Corporation.

[DOT] U.S. Department of Transportation. 1994. Federico F. Peña. World Wide Web. URL:http://www.dot.gov/dotinfo/general/sec.html.

[Flanagan] Flanagan, William A. 1994. *ATM Users Guide*. New York: Flatiron Publishing.

[FORE] FORE Systems, Inc. 1995. ForeThought Bandwidth Management Available Now. *ForeWords*. Spring.

[Forum] The ATM Forum. 1995. The ATM Forum Approves ATM LAN Emulation Specification: Users Can Preserve Existing LAN Environments While Migrating to ATM. World Wide Web. URL: http://www.atmforum.com/atmforum/LANE.html.

[Industrial] Industrial Computer Source. 1995. 1995 Communication Source-Book. San Diego, California. Industrial Computer Source.

[Keen] Keen, Peter G. W., and J. Michael Cummins. 1994. *Networks in Action, Business Choices and Telecommunications Decisions*. Belmont, California: Wadsworth.

[LANplex] 3Com Corporation. 1994. Product literature. LANplex 2000 Family, Departmental Switching Hubs. Santa Clara, California: 3Com Corporation.

[Lippis] Lippis, Nicholas J. III. 1995. Lecture entitled *Building Networks with ATM and Switching Technology: Practical Implementations*, 27 March. Networld + Interop Spring '95, Las Vegas, Nevada.

[Morency1] Morency, John. 1994. Buyer's Guide, ATM Switch Decision Can Wait. *Network World*. September 19, 70.

[Morency] Morency, John. 1995. Lecture entitled *Building networks with ATM and Switching Technology: Practical Implementations*, 28 March. Networld + Interop Spring '95, Las Vegas, Nevada.

[NBII] 3Com Corporation. 1995. Product literature. NETBuilder II High Performance Bridge/Router. Santa Clara, California: 3Com Corporation.

[Partridge] Partridge, Craig. 1994. *Gigabit Networking*. Reading, Massachusetts: Addison-Wesley Publishing Group.

[Saliba] Saliba, Waldir, Senior Network Engineer, Intermodal Data Network. 1995. Interview by author, 11 May, Washington, D.C.

[Schireson] Schireson, Max. 1994. Hubs Stretch Bandwidth. *PC Week*. 14 March, v. 11, nr. 10, N1.

[Siemens] Händel, Rainer, Manfred N. Huber, and Stefan Schröder. Siemens AG. 1994. *ATM Networks: Concepts, Protocols, Applications*. Second edition. Electronic Systems Engineering Series. Workingham, England: Addison-Wesley Publishing.

[SNCI1] Strategic Networks Consulting, Inc. 1995. The 1995 Switched Ethernet Evaluation. World Wide Web. URL:http://www.snci.com/execint.htm#intro.

[SNCI] Strategic Networks Consulting Inc. 1995. Tutorial Notes. Building Networks with ATM and Switching Technology: Practical Implementations. Softbank Exposition and Conference Company.

[Stall94] Stallings, William. 1994. *Data and Computer Communications*, 4th edition. New York: Macmillan Publishing.

[UNH] University of New Hampshire InterOperability Lab. 1994. ATM Consortium Test Suites. World Wide Web. URL:http://www.iol.unh.edu/consortiums/atm/atm_testsuite.html.

[WSJ] Hill, G. Christian, and William M. Bulkeley. 1995. 3Com is Buying Chipcom for $775 Million in Stock, Creating a Network-Gear Giant. *Wall Street Journal*, 28 July, A3.

[Xylan2] Xylan Corporation. 1995. Product literature. Total LAN Switching, Powerful ATM Access. Calabasas, California: Xylan Corporation.

[Xylan] Author's discussion with Xylan engineer at Networld + Interop Spring '95. 1995. Las Vegas, Nevada.

BIBLIOGRAPHY

Alan, Frank. 1995. Fiber Optics for Networks. Tutorial: Lesson 75. *LAN Magazine.* November, v. 9, nr. 12, 25.

American Institute. 1994. *Hands-On Migrating to ATM.* New York: American Institute.

Armbruster, H., and H. Schneider. 1990. Phasing-in the Universal Broadband ISDN: Initial Trials for Examining ATM Applications and ATM Systems. In *International Conference on Integrated Broadband Services and Networks, Conference Publication No. 329, 15-18 October, London*, by IEE.

Bellman, Robert B. 1995. An Enterprise View of ATM.. *Supplement to Business Communications Review*, April.

Bonatti, Mario, Roberta Gobbi, Giacinto Dammicco, Ugo Mocci, Massimo Lippi, Maurizio Montagna, Pietro Polese, and Giovani Roso. 1992. B-ISDN Economical Evaluation in Metropolitan Areas. *IEEE Journal on Selected Areas in Communications.* December, v. 10, Issue 9, 1489.

Broadband Publishing Corporation. 1995. The ATM Report Guide to ATM Switching Products - Part III. *The ATM Report.* 23 March, v. 3, No. 2, 1.

Buddenberg, Rex. 1995. Standards: Search for the Interoperability Keys. Lecture notes.. World Wide Web. URL:http://dubhe.cc.nps.navy.mil/~budden/lecture.notes/stds_views-_fn.html#fn1. June.

Bussey, James E. III. 1995. *Faking it with Work Content Language: Four Dimensional, High-Tech Buzzword Generator.* Washington, D.C. : Department of Transportation.

Chao, H. Jonathan, Dipak Ghosal, Debanjan Saha, and Satish K. Tripathi. 1994. IP on ATM Local Area Networks. *IEEE Communications Magazine*, August, v. 32, Issue 8, 52.

Chernicoff, David P. 1994. The Real Story of ATM; Moving to an Enterprise Topology. *PC Week.* 16 May, v. 11, nr. 19, N1.

Codenoll Technology Corporation. 1993. *The Fiber Optic LAN Handbook.* 5th Edition. Yonkers, New York: Codenoll Technology Corporation.

Cuthbert, L.G. 1991. Issues in the Planning of Broadband ATM Networks. In *Third IEE Conference on Telecommunications, Conference Publication No. 331, 17-20 March, Edinburgh,* by IEE. London: IEE, 334.

Dahod, Ash. 1994. Available Today, LAN Switching Meets Demand for Higher LAN Bandwidth. *Special Supplement: Computing Canada Report on Communications and Networking.* 14 September, v.20, nr. 19, S10.

Durr, Michael. 1994. Moving Step by Step toward ATM; Asynchronous Transfer Mode is the Network of the Future, but, When ? *LAN Times*, 11 April, v. 11, nr. 7, 47.

Feibel, Werner. 1995. *Novell's Complete Encyclopedia of Networking.* Alameda, California: Sybex.

Fischer, Wolfgang, Eugen Wallmeier, Thomas Worster, Simon P. Davis, and Andrew Hayter. 1994. Data Communications Using ATM: Architectures, protocols, and Resource Management. *IEEE Communications Magazine*, August, v. 32, Issue 8, 24.

Freeman, Roger L. 1991. *Telecommunication Transmission Handbook.* 3rd Edition. Wiley Series in Telecommunications, ed. Donald L. Schilling, New York: John Wiley & Sons.

Gerla, Mario, Tsung-Yuan Tai, and Giorgio Gallassi. 1992. LAN/MAN Interconnection to ATM: A Simulation Study. In *Eleventh Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM '92, 4-8 May, Florence,* by the Institute of Electrical and Electronic Engineers. New York: IEEE, 2270.

Giancarlo, Charlie. 1995. KALPANA: A Perspective from the LAN Switch Pioneers. *Supplement to Business Communications Review*, April.

Humphrey, Marlis. 1995. ATM Forum Update: 1995 Issues and Goals. *Telecommunications*. June, 40.

Jiandani, Soni. 1995. Migrating to ATM — A Phased Strategy based on Switching Technology. *Supplement to Business Communications Review*, April.

Khalil, Khalid M., and Paul A. Spencer. A Systematic Approach for Planning, Tuning, and Upgrading Local Area Networks. In *Conference Record. IEEE Global Telecommunications Conference, GLOBECOM '91,* by the Institute of Electrical and Electronic Engineers. New York: IEEE, 658.

Linke, Scott. 1994. ATM Physical Interfaces. In *Seminar Handbook, B-ISDN/ATM Testing Technologies Seminar 1994,* by Hewlett-Packard Company. U.S.A.: Hewlett-Packard Company.

Lippis, Nicholas J. III. 1993. ATM LANs: Software Will Make the Difference. *Data Communications.* August, 29.

Lippis, Nicholas J. III. 1993. ATM to the Desktop? I Don't Think So. *PC Magazine.* December 21, NE45.

_____. and John Morency. 1994. On the Road to ATM. *PC Magazine.* January 25, NE1.

_____. 1992. ATM's Biggest Benefit: Better Net Management. *Data Communications.* October, 27.

Macaraeg, Brendon. 1994. Switching Hubs Free More Network Bandwidth...Today. *PC Magazine.* 28 June, v. 13, nr. 12, 32.

Mandeville, Robert. 1995. The ATM Stress Test: Which Switches Survived? *Data Communications,* March.

Mazzola, Mario. 1995. Switching Technologies Take Off. *Supplement to Business Communications Review,* April.

Mocci, U., and G. Dammicco. 1990. Evolutionary Strategies of Interoffice Networks. In *International Conference on Integrated Broadband Services and Networks, Conference Publication No. 329, 15-18 October, London,* by IEE. London: IEE, 23.

ONET Data Communication Technologies. 1995. MetaSwitching™: Next-Generation Ethernet Switching from ONET Data Communication Technologies. ONET Position Paper. Cambridge, Massachusetts: ONET Data Communication Technologies.

Pitts, J. M. and L.G. Cuthbert. 1993. ATM Network Dimensioning: An Overview. In *Fourth IEE Conference on Telecommunications, Conference Publication No. 371, 18-21 April, Manchester, U.K.,* by IEE. London: IEE, 269.

Powers, John T. And Henry H. Stair II. 1990. *Megabit Data Communications: A Guide for Professionals.* Englewood Cliffs, New Jersey: Prentice-Hall.

Rowe, Stanford H. 1991. *Business Telecommunications.* 2nd Edition. New York: Macmillan Publishing.

Ruiu, Dragos. 1994. ATM Internetworking. In *Seminar Handbook, B-ISDN/ATM Testing Technologies Seminar 1994,* by Hewlett-Packard Company. U.S.A.: Hewlett-Packard Company.

Ruiu, Dragos. 1994. B-ISDN Overview. In *Seminar Handbook, B-ISDN/ATM Testing Technologies Seminar 1994,* by Hewlett-Packard Company. U.S.A.: Hewlett-Packard Company.

Schatt, Stan. 1993. *Understanding Local Area Networks.* 4th Edition. New York: Sams Publishing.

Scott, Andrew. 1994. ATM Signaling. In *Seminar Handbook, B-ISDN/ATM Testing Technologies Seminar 1994,* by Hewlett-Packard Company. U.S.A.: Hewlett-Packard Company.

Stassinopoulos, G. I., I. S. Venieris, and R. Carli. 1990. ATM Adaptation Layer Protocols and IEEE LAN Interconnection. In *Proceedings. 15th Conference on Local Computer Networks,* by IEEE Computer Society. Los Alamitos, California: IEEE Computer Society.

The ATM Forum. *Introduction to ATM.* Foster City, California: The ATM Forum.

Walker, Julie. 1994. New Performance Challenges. In *Seminar Handbook, B-ISDN/ATM Testing Technologies Seminar 1994,* by Hewlett-Packard Company. U.S.A.: Hewlett-Packard Company.

INITIAL DISTRIBUTION LIST

		No. Copies
1.	Defense Technical Information Center Cameron Station Alexandria, Virginia 22304-6145	2
2.	Library, Code 013 Naval Postgraduate School Monterey, California 93943-5002	2
3.	LCDR James W. Kelly Maintenance and Logistics Command Pacific (ts) Coast Guard Island, Building 54a Alameda, CA 94501-5100	2
4.	Director, Intermodal Data Network U.S. Department of Transportation Washington, DC 20593	2
5.	Commandant (G-TTM) U.S. Coast Guard 2100 Second Street, S.W. Washington, DC 20593-0001	1
6.	Professor Myung Suh Department of Systems Management, Code SM/Su Naval Postgraduate School Monterey, CA 93943-5103	1
7.	Professor Rex Buddenberg Department of Systems Management, Code SM/Bu Naval Postgraduate School Monterey, CA 93943-5103	1
8.	Commandant (G-L) U.S. Coast Guard 2100 Second Street, S.W. Washington, DC 20593-0001	1

9. CAPT Robert Williams, USMC 1
Department of Systems Management, Code 36
Naval Postgraduate School
Monterey, CA 93943-5103

10. LT H. S. Li, R.O.C. on Taiwan Navy 1
Department of Systems Management, Code 36
Naval Postgraduate School
Monterey, CA 93943-5103